Summary of FE-I Design Status

K. Einsweiler, LBNL

Critical Issues from June FE-I Review

- Perforamance of the analog front-end observed in test chips
- Front-end configuration, shielding, and power distribution for analog front-end

Progress in understanding analog front-end performance

- Detailed comparisons of TSMC and IBM chip performance with HSPICE

Present Status

- Layout and Schematics
- Overall Verification

Analog Issues from the Review

Threshold dispersion and matching:

- Have carried out several studies of impact of VT matching on present front-end
- First simulations just used a voltage source to shift the VT of each critical transistor in the preamplifier and second stage by a known amount, and then scaled to the expected VT mis-match for that device. Unfortunate fact is that most transistors in the preamp and second amplifier contribute to the threshold dispersion.
- Two analyses has been done. One replaces each FET by a subcircuit, and varies the VT in accordance with W and L, but only uses DC operating points to estimate the dispersion. The second approach actually does a threshold scan for each set of VT shifts, and then computes a dispersion. The two analyses agree well.
- Each device had its VT modified using sigma taken from the thesis of G. Anelli:

Standard PMOS VT matching versus device size

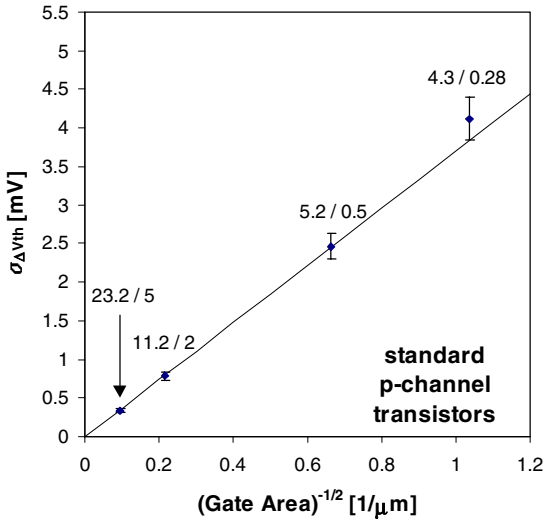

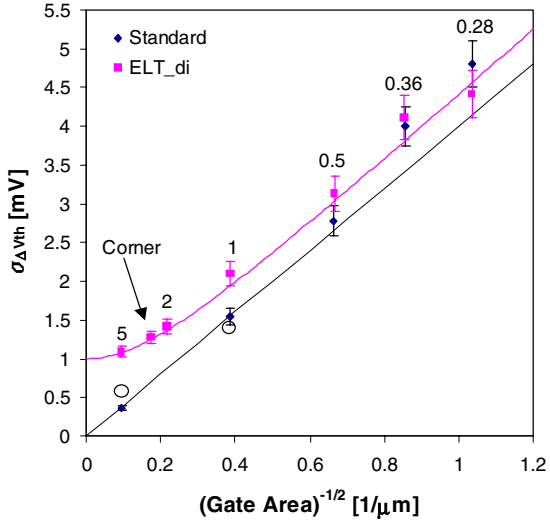

Enclosed NMOS VT matching versus device size

- Threshold simulation, moving each device by ΔVT = +5mV:

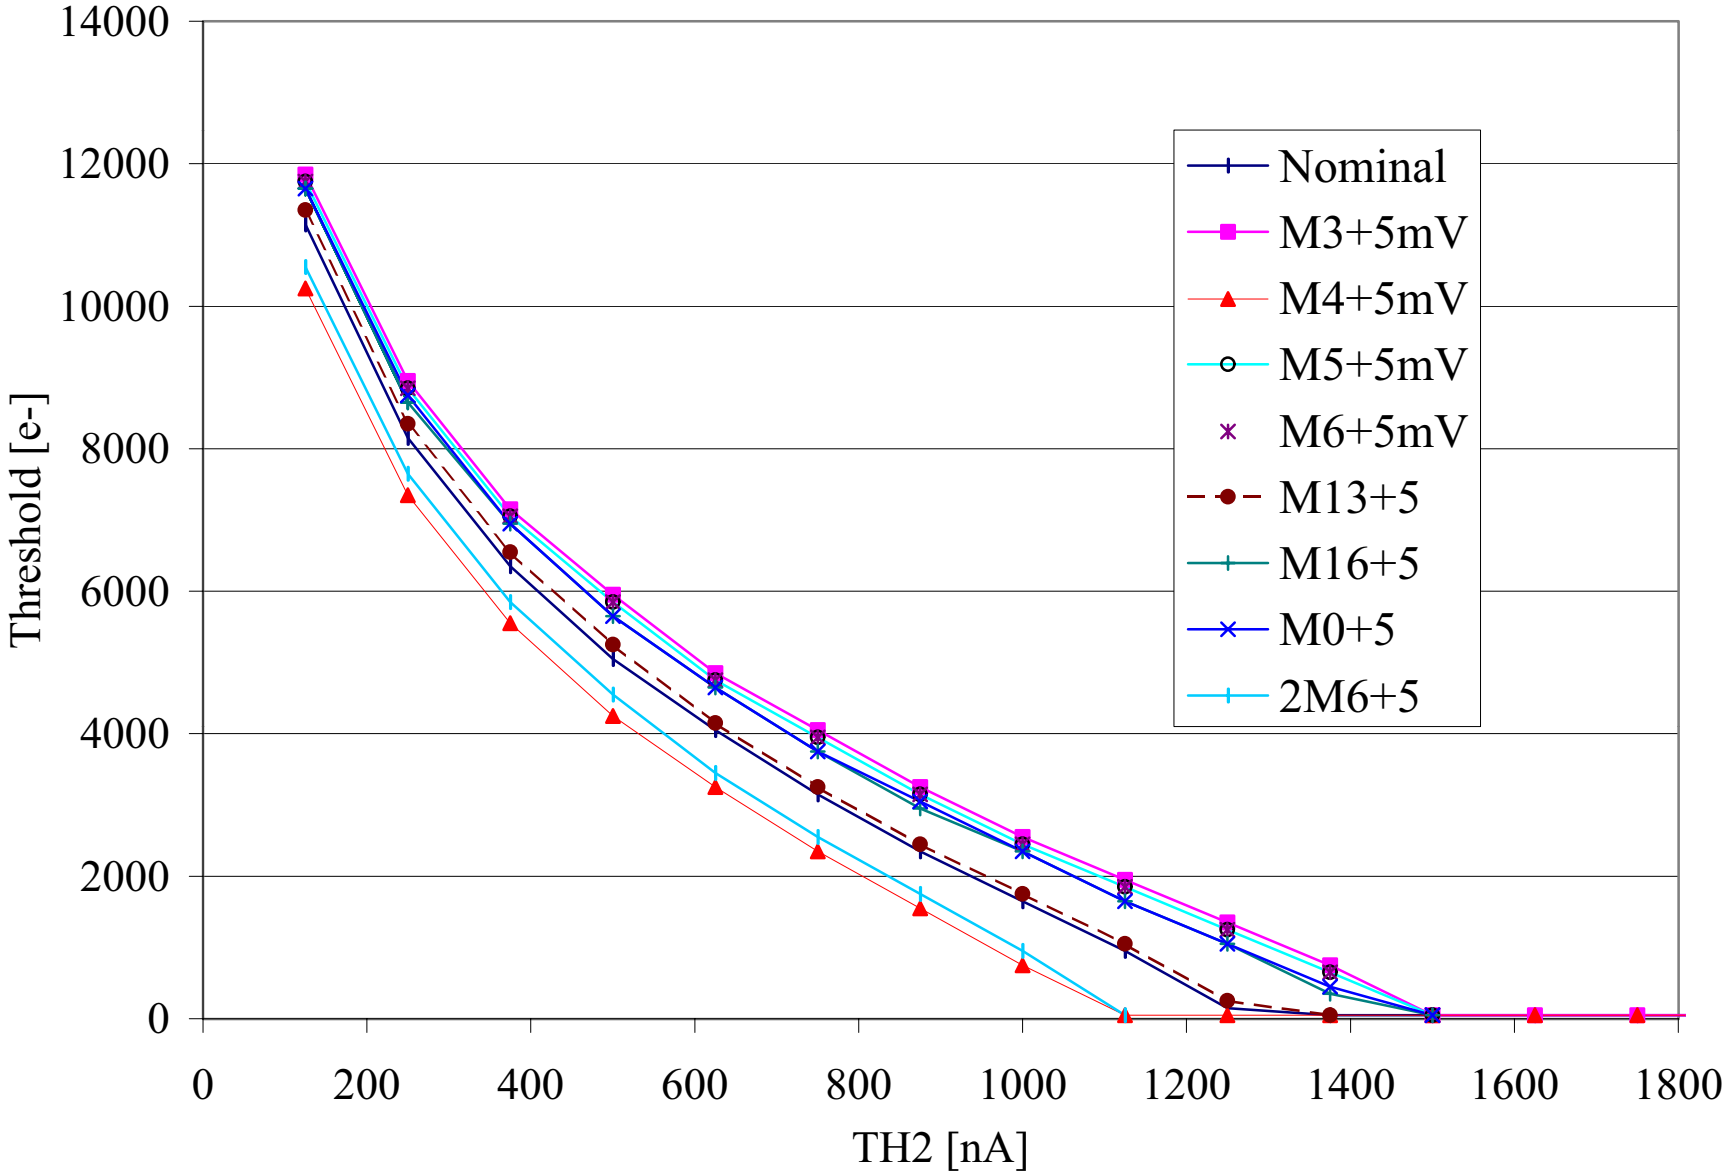

- Need to combine contributions in quadrature, but behavior looks like measurement.

- Results of Monte Carlo simulation, sliced at one value of ITH2 (threshold):

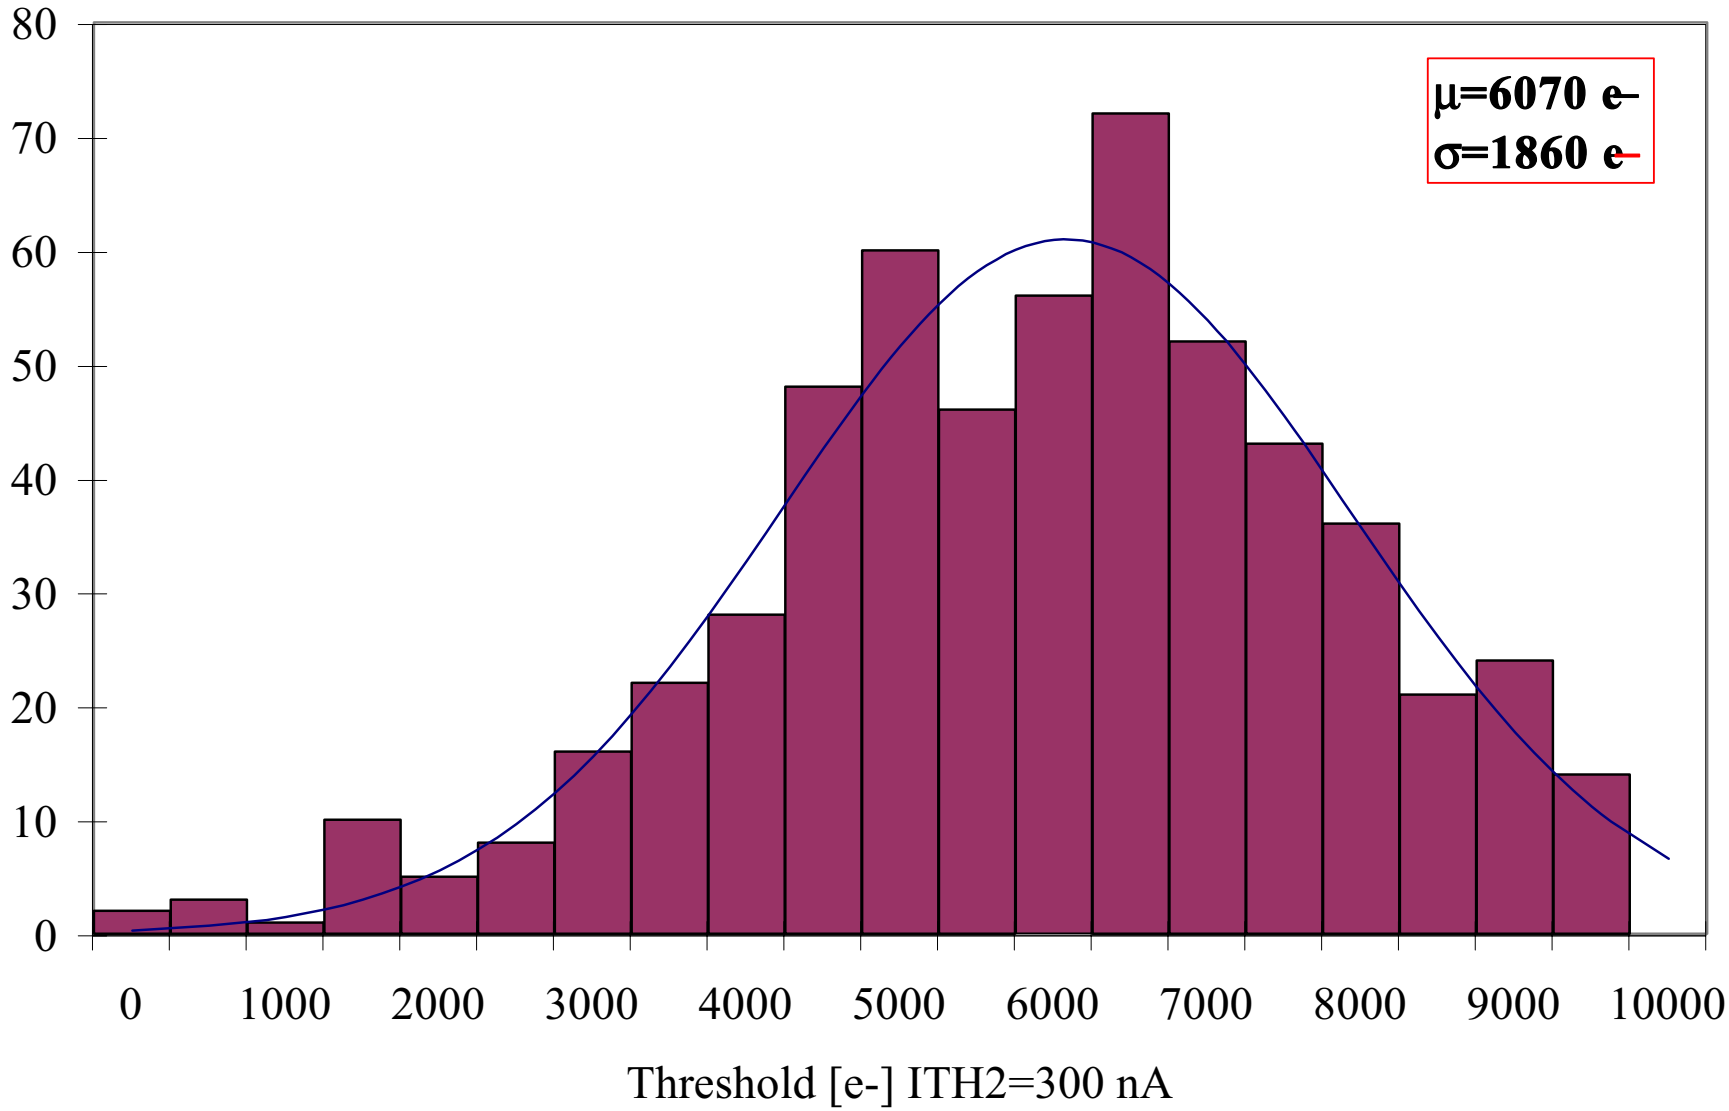

- Have not included sources of chip-chip variations, but matching contributions are consistent with the magnitude of the observed threshold dispersion.

Threshold dispersion before and after tuning in IBM test chip:

- Example of observed raw dispersion and noise from threshold scan:

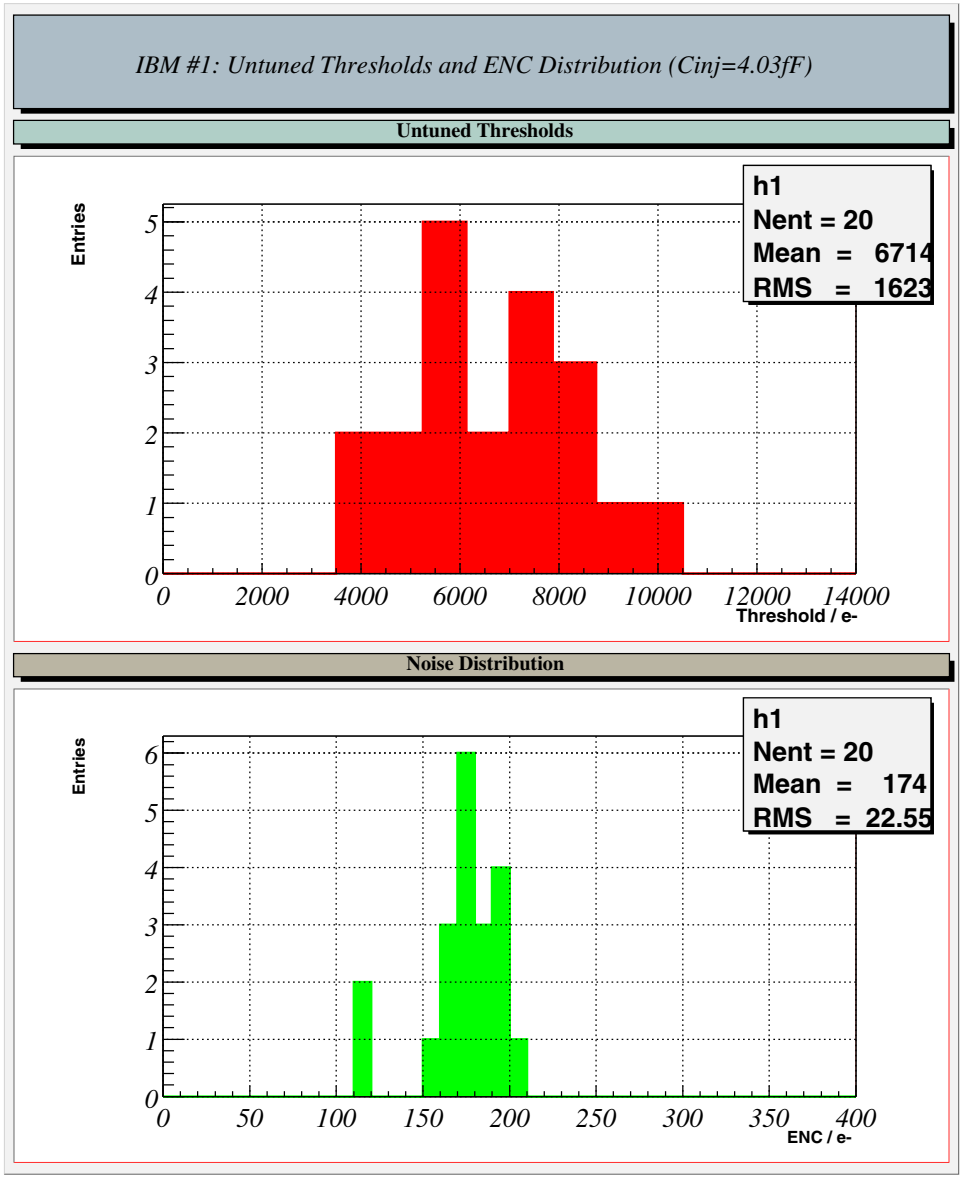

Observed dispersion is consistent with estimate from matching Monte Carlo, considering only VT matching of devices in preamp and second amplifier.

This result uses the latest and best estimate for the injection capacitor.

- Tuning of threshold using 5-bit pixel TDACs:

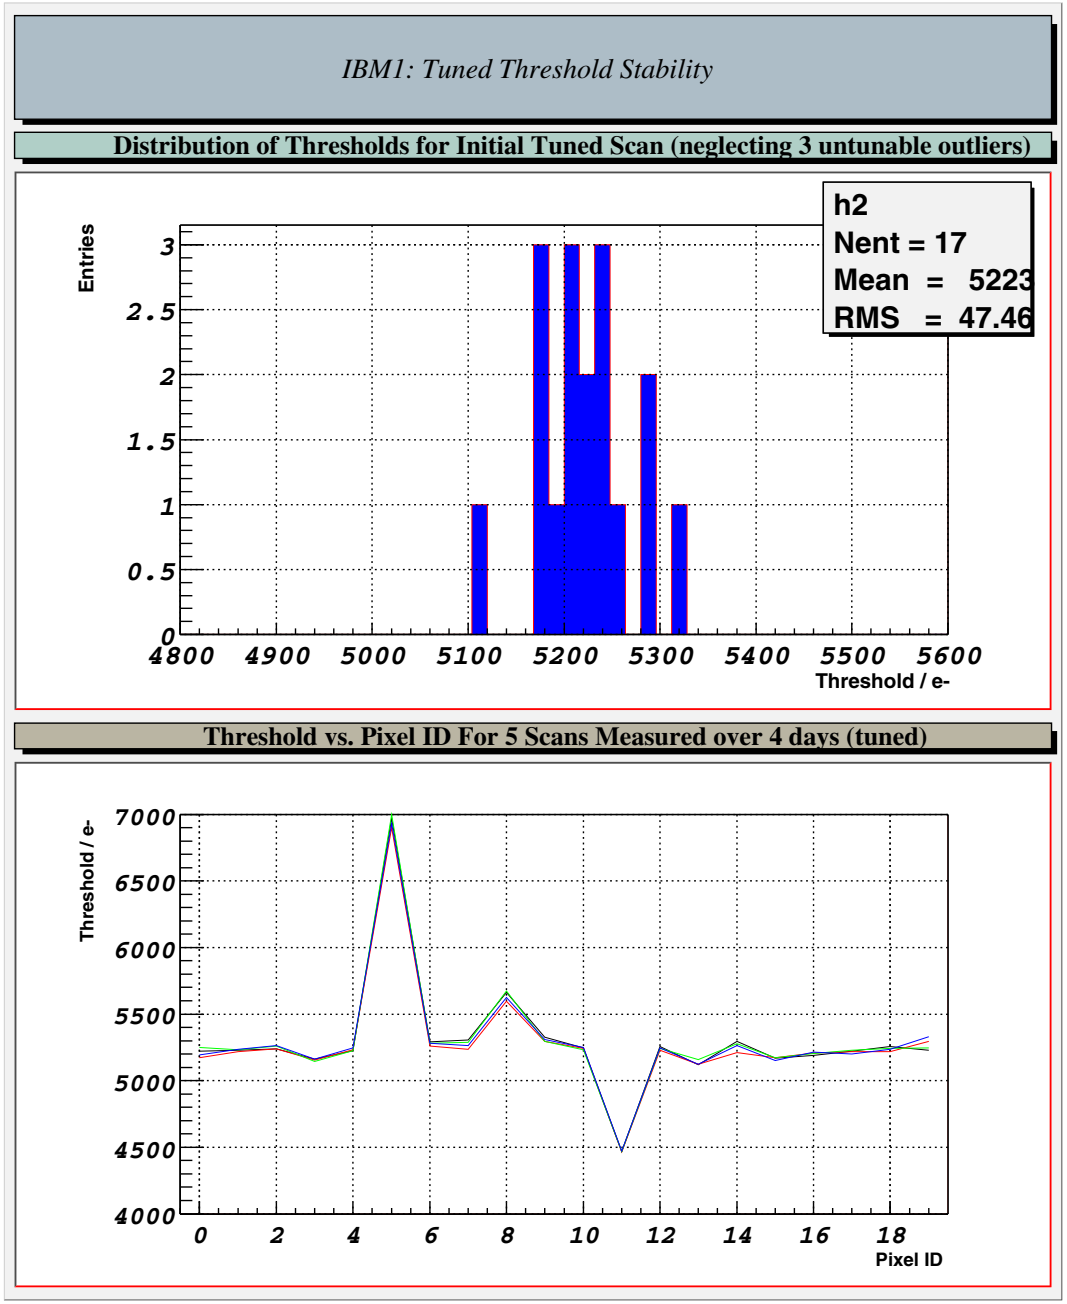

Range of TDAC is as expected in the test chip, but it was never intended to work with such large threshold dispersions. Therefore, in the present test chips, not all channels can be tuned to a common threshold point.

Initial dispersion of about 1600e was successfully tuned down to an RMS of about 50e, for those 17 channels which were within range of the TDAC.

Stability of the tuning is plotted here over a 4 day interval (three poorly tuned channels are apparent). It was very stable over a period of more than one month of testing.

How are we responding to this dispersion for this run ?

- Major difference between this design and all earlier designs is the very low gain of the preamplifier. Previously, we used a feedback capacitor in the range of 3-4fF and a single-stage front-end. In the test chips, we have a drawn feedback capacitance of Cfb=14fF. When considering the Cfb which is relevant for the preamp gain, there is a significant additional contribution from the parasitic capacitance of the feedback transistors, making the total around 20fF. This makes the preamp gain about 5 times lower than in earlier designs.
- Analysis showed that many devices were responsible for dispersion. The dispersion in the second amplifier was larger than in the preamplifier, so a simple AC-coupling approach would not have helped significantly (30-40% improvement). Many transistors whose size could be increased for better matching must also be small to avoid large parasitic capacitances on critical nodes, so no easy fixes...
- We believe that a significant improvement in the threshold dispersion of the design would require major changes to the threshold control, and would require another test chip iteration to reduce the risk to an acceptable level. We intend to proceed with this work, but not until the present chip is submitted for an engineering run.
- We therefore considered changes that are rather modest, and that we believe are small extrapolations from the present design. The path chosen for this run is to reduce the value of the feedback capacitance. One design would have the nominal 10fF feedback, and the second a more aggressive 5fF feedback. The 5fF chip is expected to have a reduction by a factor of two in the threshold dispersion.

Power distribution, shielding, power supply rejection:

- Present preamplifier design is single-ended with a PMOS input device and a folded cascode. This design has considerable sensitivity to noise on the power supply, since this modulates Vgs of the input transistor in roughly the same way as a true input at the gate does. The ground node has little noise sensitivity.
- Our strategy for power distribution within the chip is consistent with this “feature”. There are two analog supply nets, VDDA and VDDREF. VDDA provides power to everything except the preamplifier. Note that everything except the preamplifier is differential, and has very good power supply rejection. Also note that in normal operation, there is very little AC current on these two nets (all inverters are on the digital supply). These two supply nets are connected by wire-bonds off-chip.
- We have also implemented significant amounts of local decoupling on the digital power supply, placing a capacitance of about 5.5pF inside of each pixel. This should minimize any spikes on the substrate and the digital supply/return nets.
- The present shielding arrangement is similar to what was discussed in the June Review. In the front-end region, M5 is VDDREF and M4 is a mixture of AGND and VDDA. The M3 layer contains many vertical bias connections, etc. Over the digital readout region of the pixel (back of the pixel), M5 remains VDDREF, and M4 is VDDA. The M3 layer contains many vertical bus connections for the pixel readout. All of these M3 signals are differential, and the asynchronous data readout lines are also reduced swing (VDD/2). This seemed to be the best compromise between shielding and power distribution (no unused shield layer was possible).

Detailed studies of Analog Test Chip performance

- Fabricated 20-pixel analog test chips in both TSMC and IBM 0.25μ processes.
- Detailed study has revealed minor and major process differences, some of which have a large effect on our design performance.
- In general, all inter-layer dielectrics are thinner in IBM, making for larger parasitic capacitances.

Impact of large PMOS well-substrate capacitance in IBM:

- In addition, we only recently realized that the IBM epi layer is quite highly doped, leading to a very large well-substrate parasitic (ten times larger than for TSMC) of almost 1fF/μ^2.
- This affects our design in three different places. First, the preamp feedback transistors used a body connection to the PMOS source to increase the dynamic range, resulting in an estimated parasitic of 60fF, which significantly degrades the risetime of our preamplifier in the IBM process. Second, in the input stage to the 50Ω buffer we use to monitor internal waveforms, a combined MUX and source-follower also used PMOS with bodies connected to their sources, adding a 130fF parasitic which degraded the risetime of the buffer by a factor of 2 (this can be compensated by doubling the source follower bias). Third, the load capacitors that we used to simulate the sensors were PMOS transistors in inversion. Their gate capacitance was lower than expected (about 60% of Cox), but in IBM, the large well capacitance doubled their value compared to TSMC.

- Comparison of preamp risetime, with/without well parasitics, as seen by 50Ω buffer:

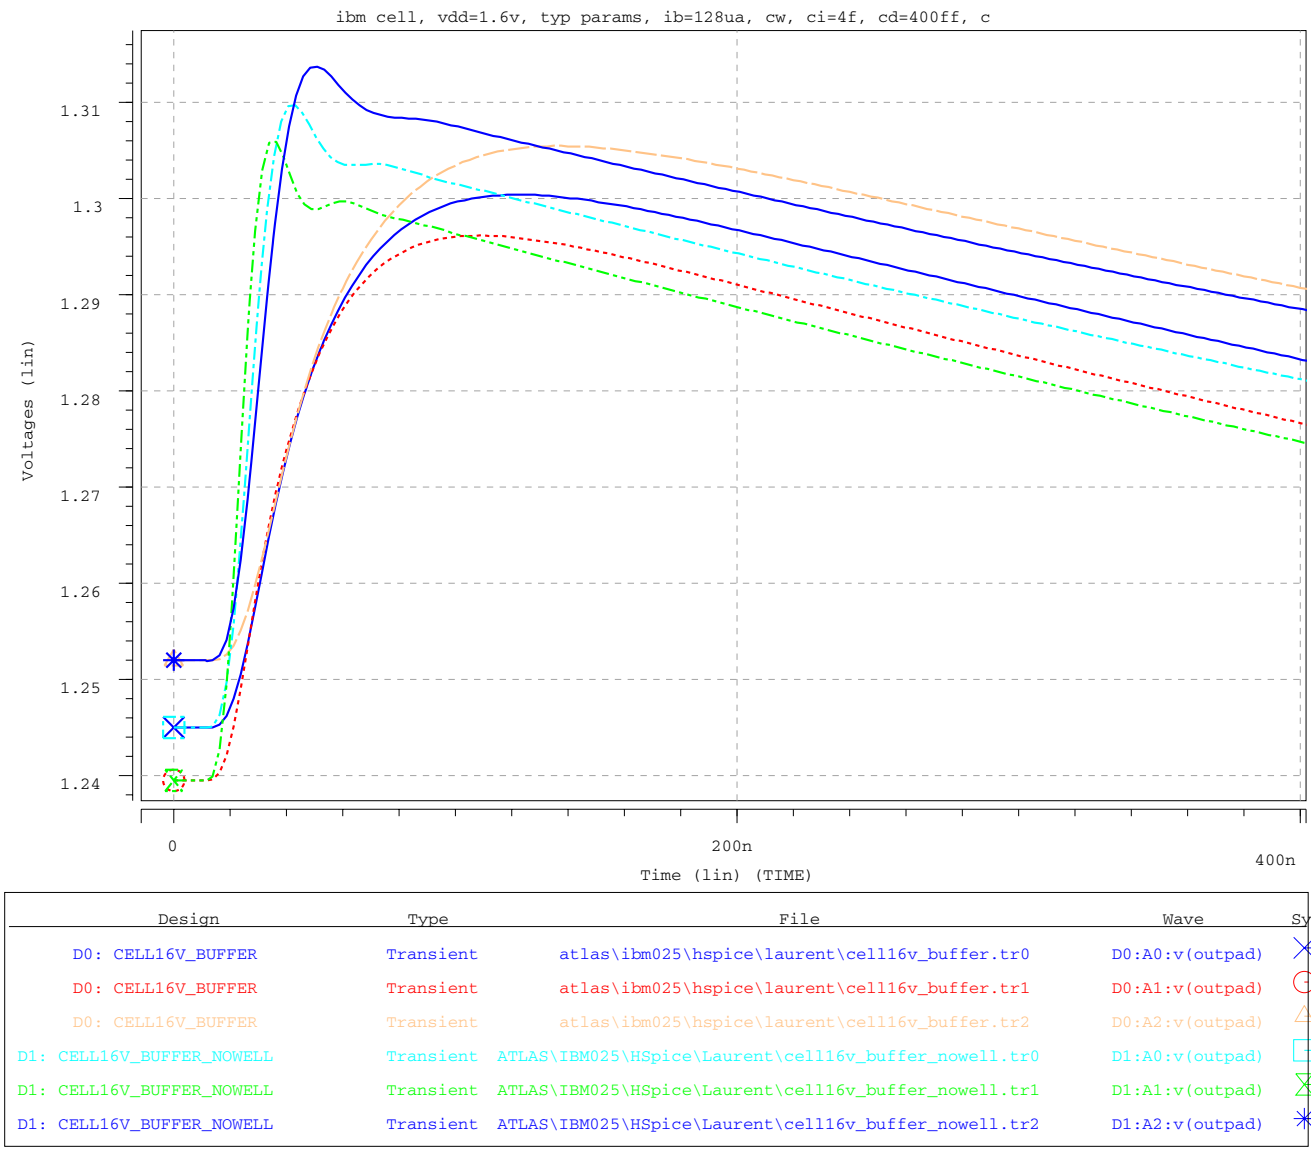

- Risetime at 10Ke with expected 400fF load degrades from 15ns to 50ns !!!

- In addition, there were some minor errors in internal signal and power routing that added parasitic capacitances and series resistances in some areas.
- We have attempted to include all of these effects in our simulations. We have tried to set up the lab measurements with equivalent parameters, with all bias currents at their nominal values.
- Some of these currents can be directly checked. The preamp bias of 8μA can be checked by measuring the change in analog power versus DAC setting, and is typically accurate to about 5%. The feedback current is mirrored out of the pixel and participates in a programmable OR chain, which allows us to directly measure the feedback current of any pixel. A standard value of IF=1nA corresponds to a feedback discharge current of 2nA.
- We have then carried out detailed comparisons of preamp gain, preamp risetime, TOT behavior, second stage gain, noise, and timewalk. The agreement is not perfect, but it is quite acceptable, typically at the 10-20% level. This has convinced us that we understand the basic performance of the front-end fairly well, and that HSPICE simulations can be used to predict the performance of any modifications that we make.

- Example comparison of preamp output amplitude (including 50Ω buffer):

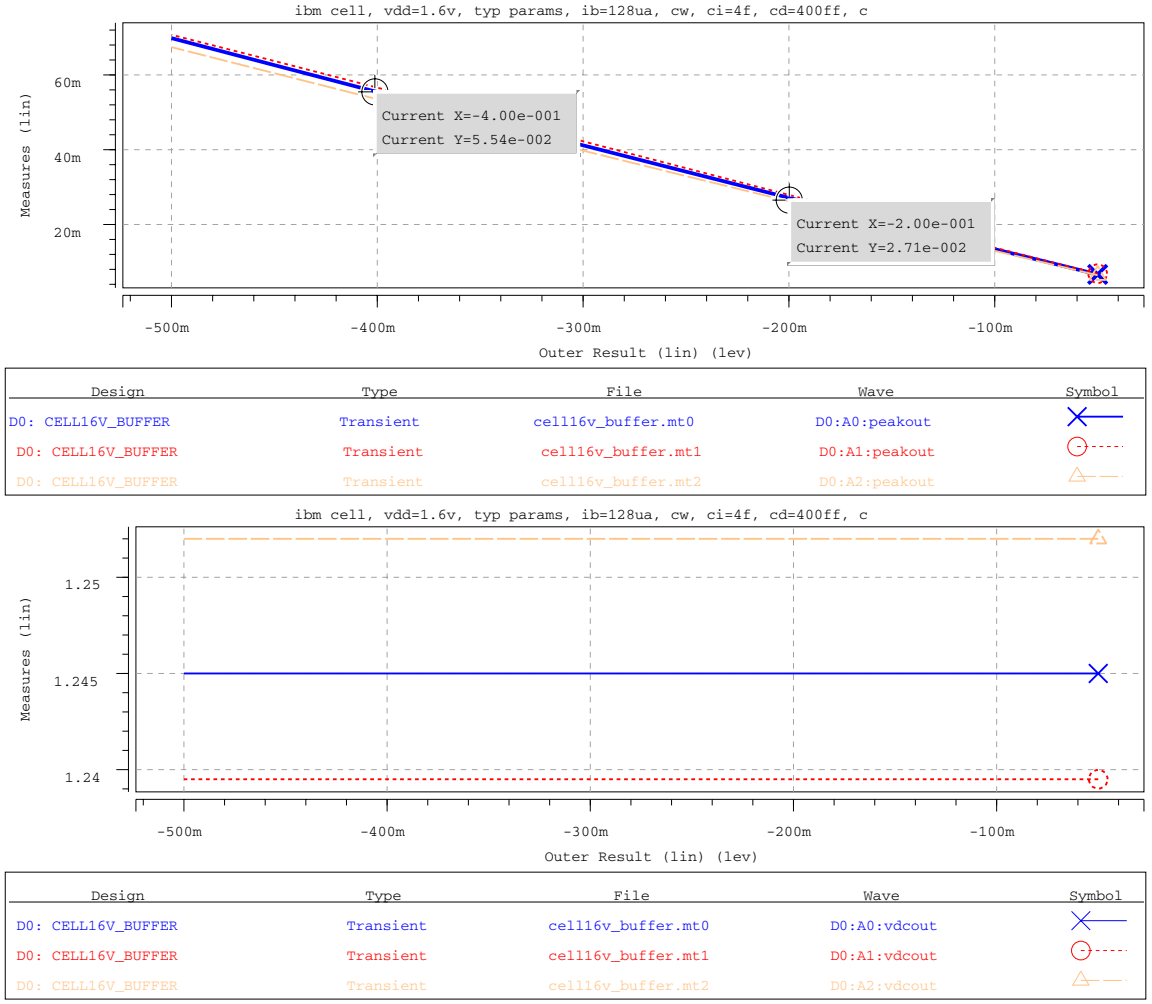

- Lab measurements for two points (roughly 5Ke and 10Ke) are 22mV and 51mV. The preamp gain and buffer gain seem to be fairly well modeled.

- Example of preamp risetime versus injected charge:

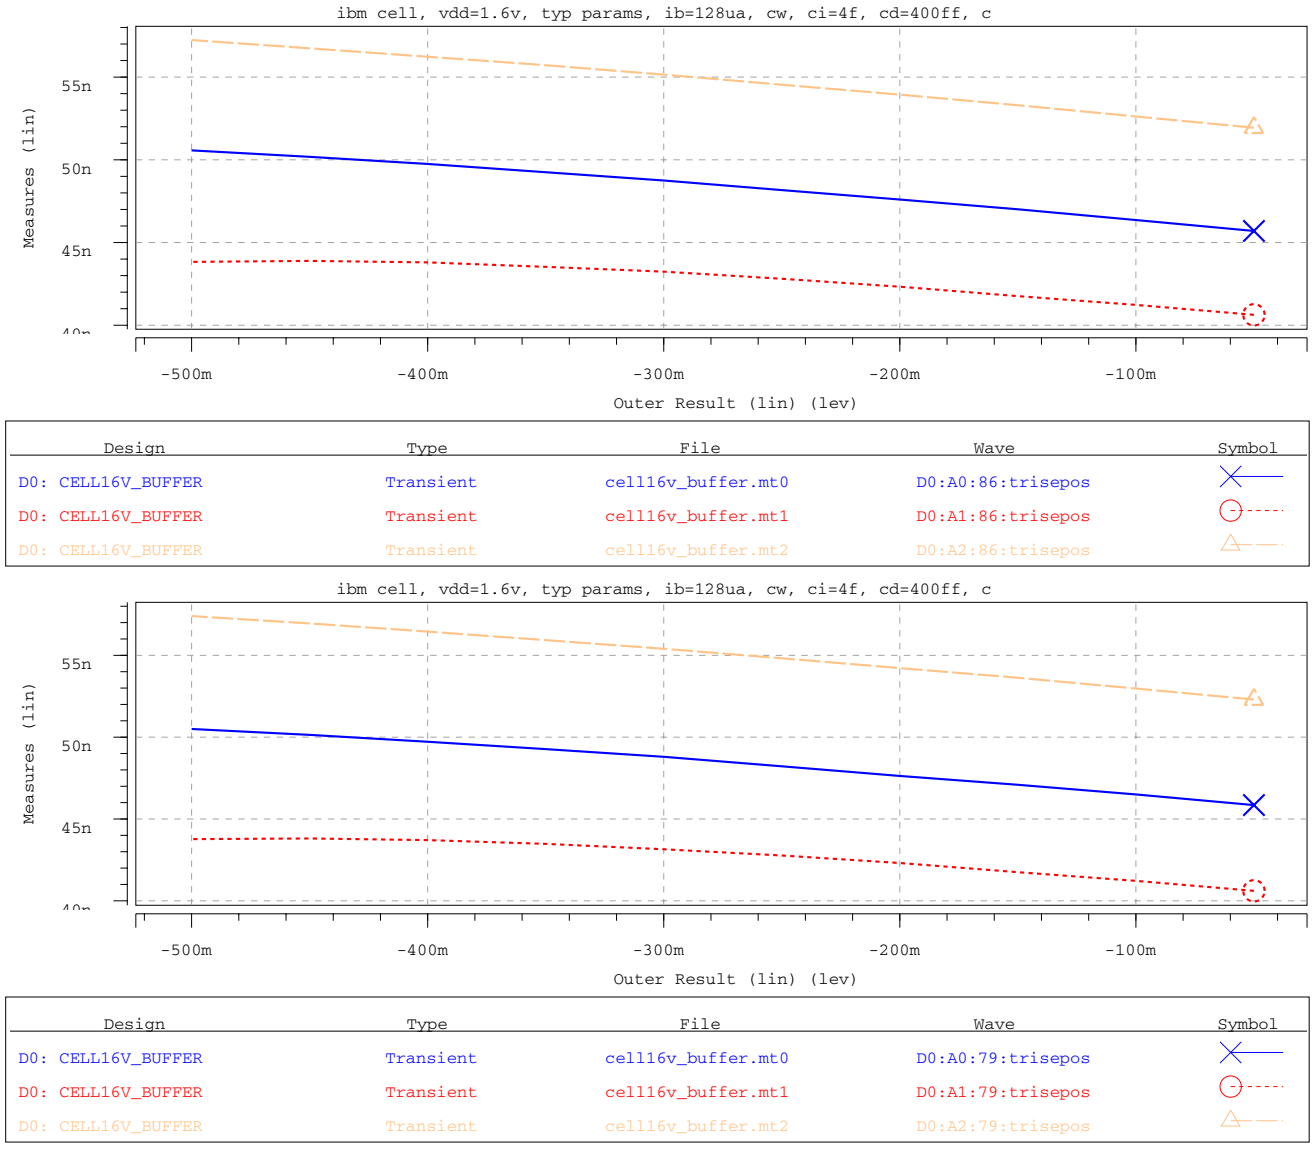

- Lower plot is after 50Ω buffer. Measurements give risetimes varying from 45ns at small amplitude to 50-55ns at larger amplitude, in good agreement.

- Example of TOT versus injected charge for IF=1nA:

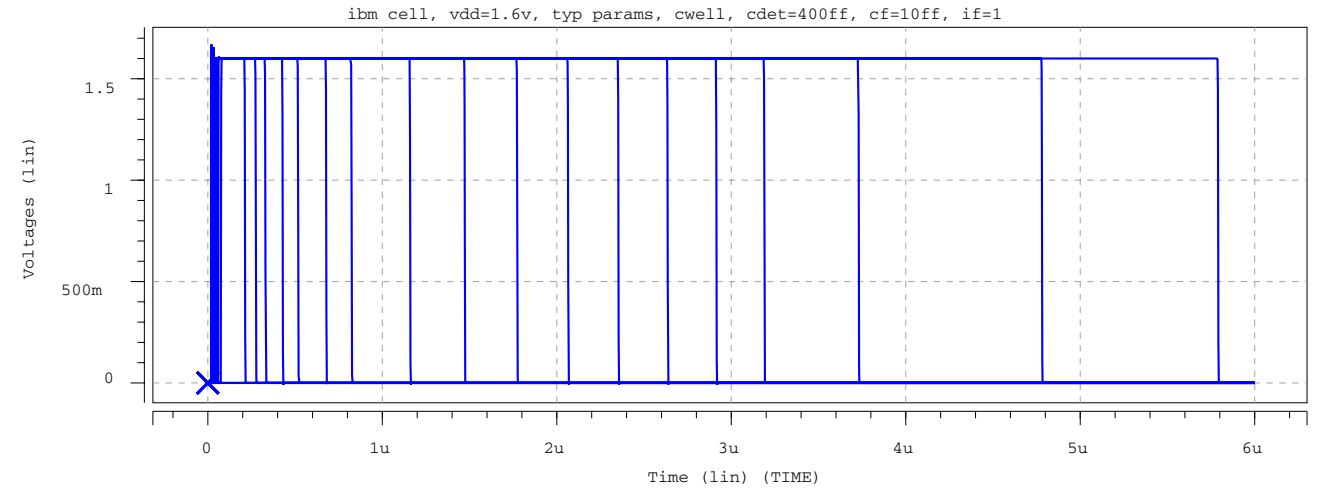

Design	Type	Wave	Symbol
D0: CELL16V_TOT	Transient	D0:A0:v(out)	

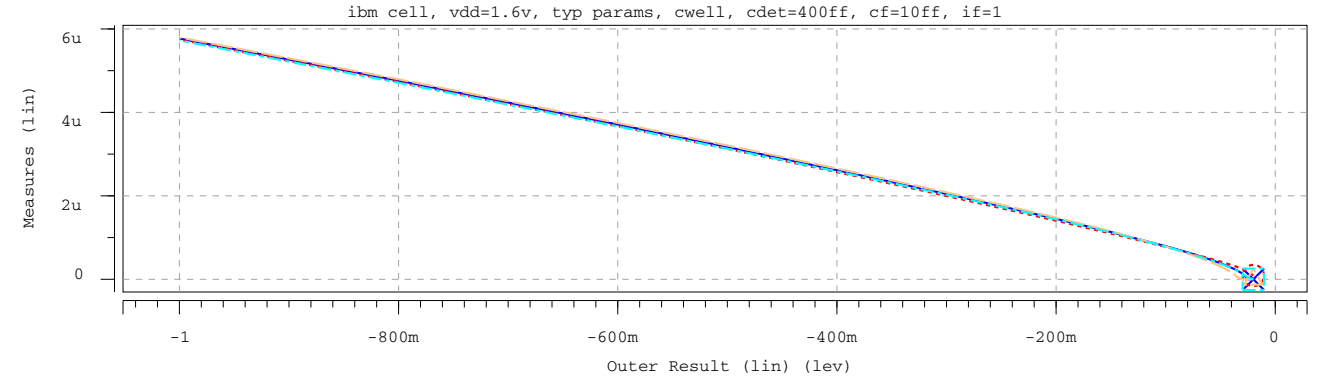

Design	Type	File	Wave	Symbol
D0: CELL16V_TOT	Transient	cell16v_tot.mt0	D0:A0:tot	
D0: CELL16V_TOT	Transient	cell16v_tot.mt1	D0:A1:tot	
D0: CELL16V_TOT	Transient	cell16v_tot.mt2	D0:A2:tot	
D0: CELL16V_TOT	Transient	cell16v_tot.mt3	D0:A3:tot	

- Very linear TOT performance observed both in the lab and in simulation. For 20Ke, measure 1.1µs in lab and about 1.3µs in simulation.

- Example of noise measurement:

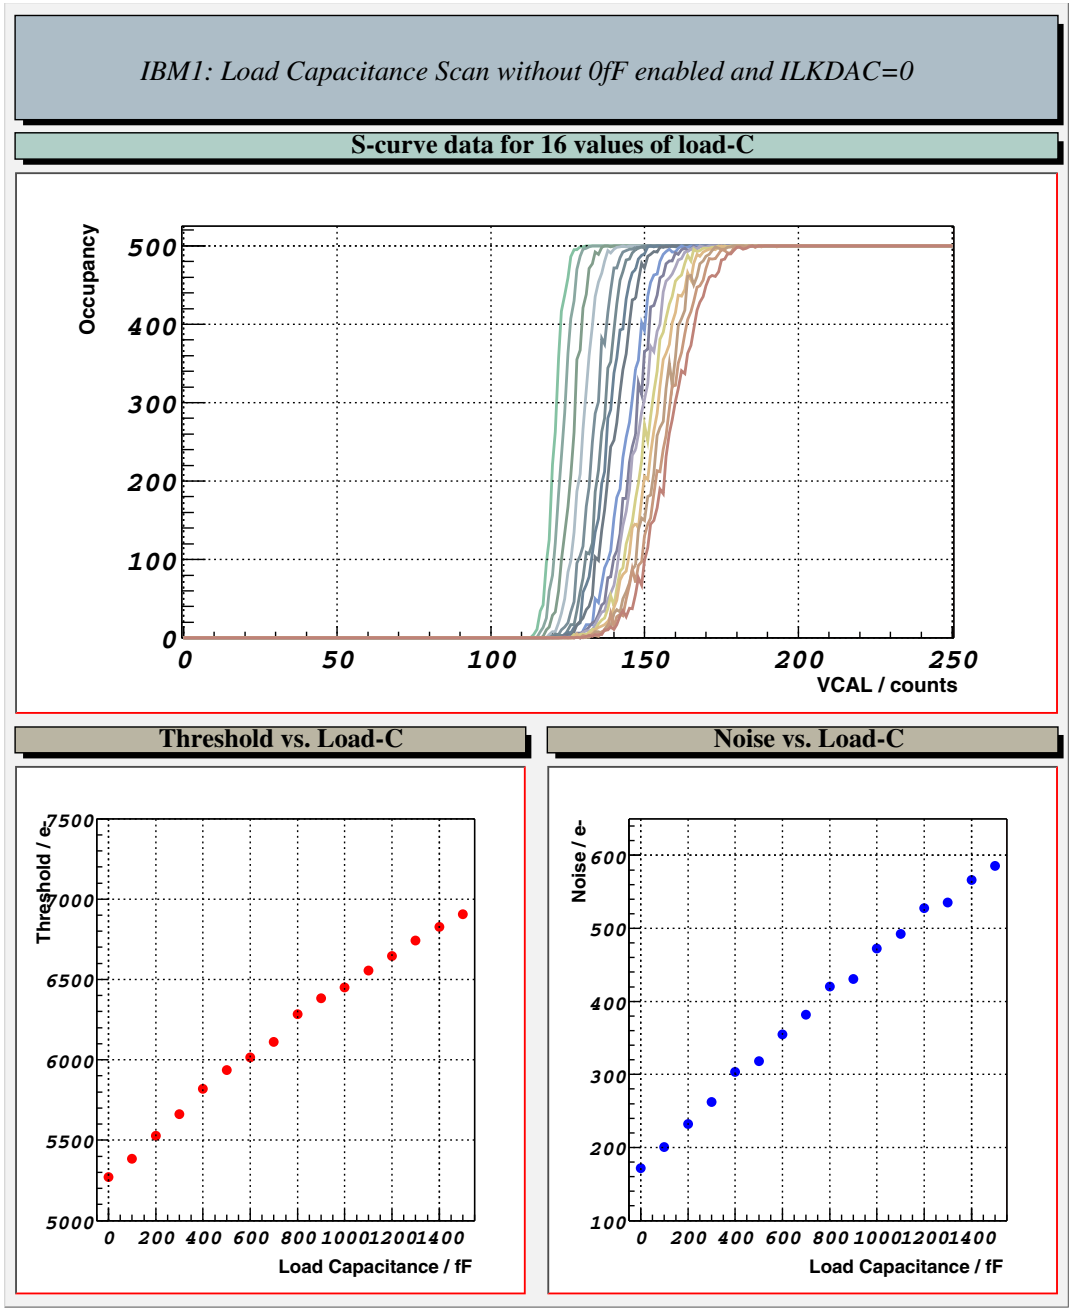

Noise and threshold are measured with different load capacitances. Including well parasitic in IBM, estimate that 100fF load is really 120fF, so results should be scaled.

These scans used an older estimate for the injection cap, and so should be scaled up by about 10%.

Leakage current can also be added. At leakage values of about 25nA expected at the end of the sensor lifetime, the leakage (parallel) and capacitive load (series) noise are about the same, so the total noise increase by about 50%.

- Example of measurement and tuning of feedback current:

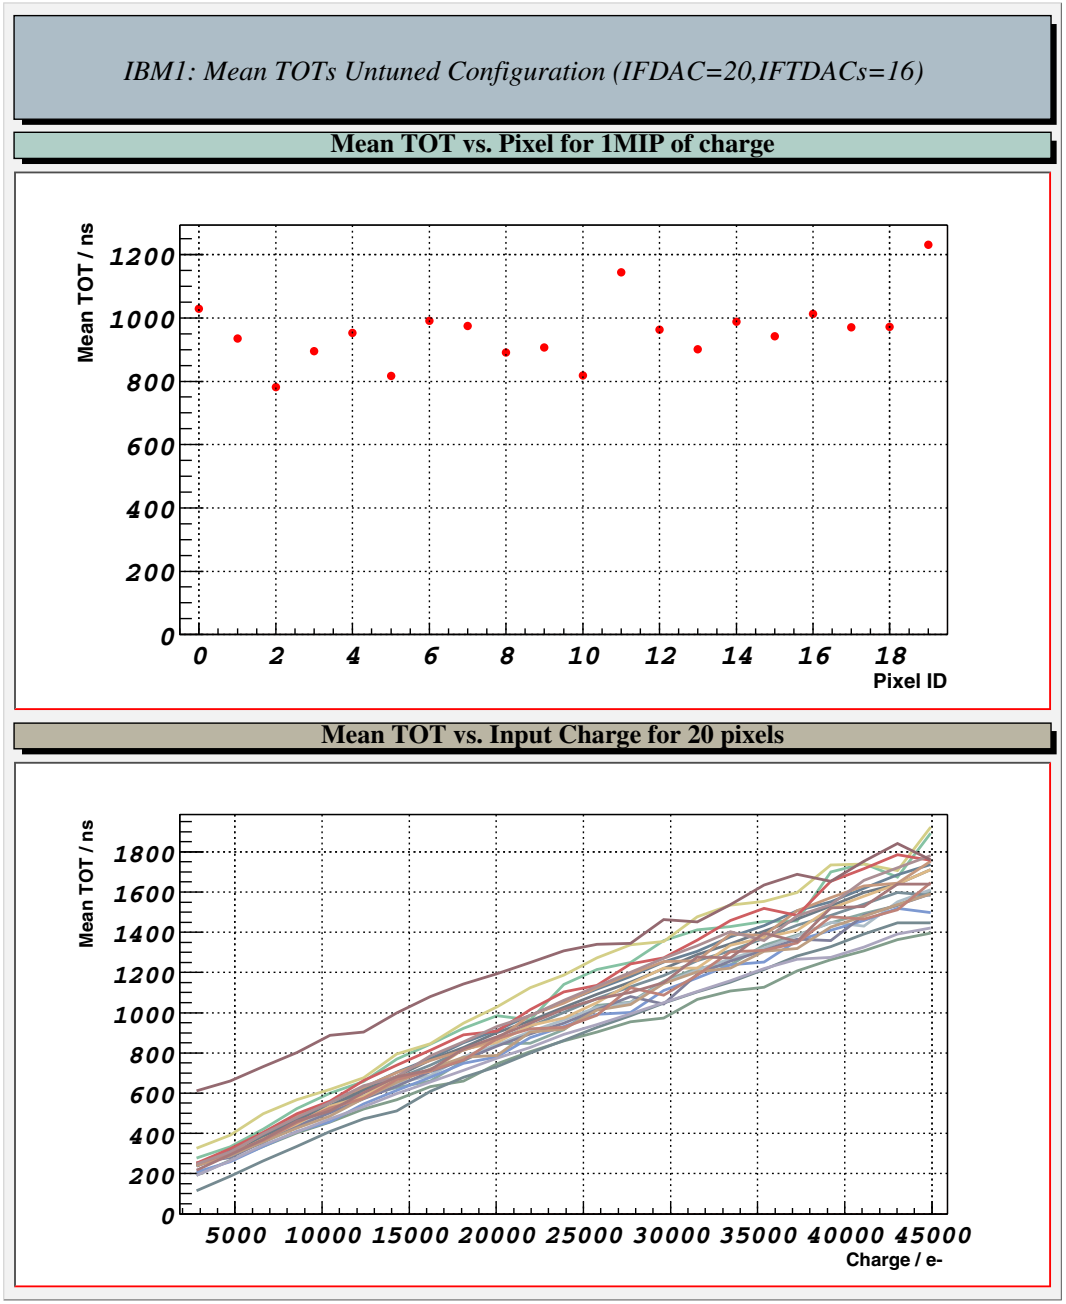

Plots here show TOT performance before tuning.

The matching of TOT depends on the matching of IF and Cfb between pixels, and this appears to be very good.

The one strange channel is the one with the 50Ω buffer and probe pad monitoring, and should be ignored.

- Performance observed after tuning IF using IFTrim DAC in each pixel:

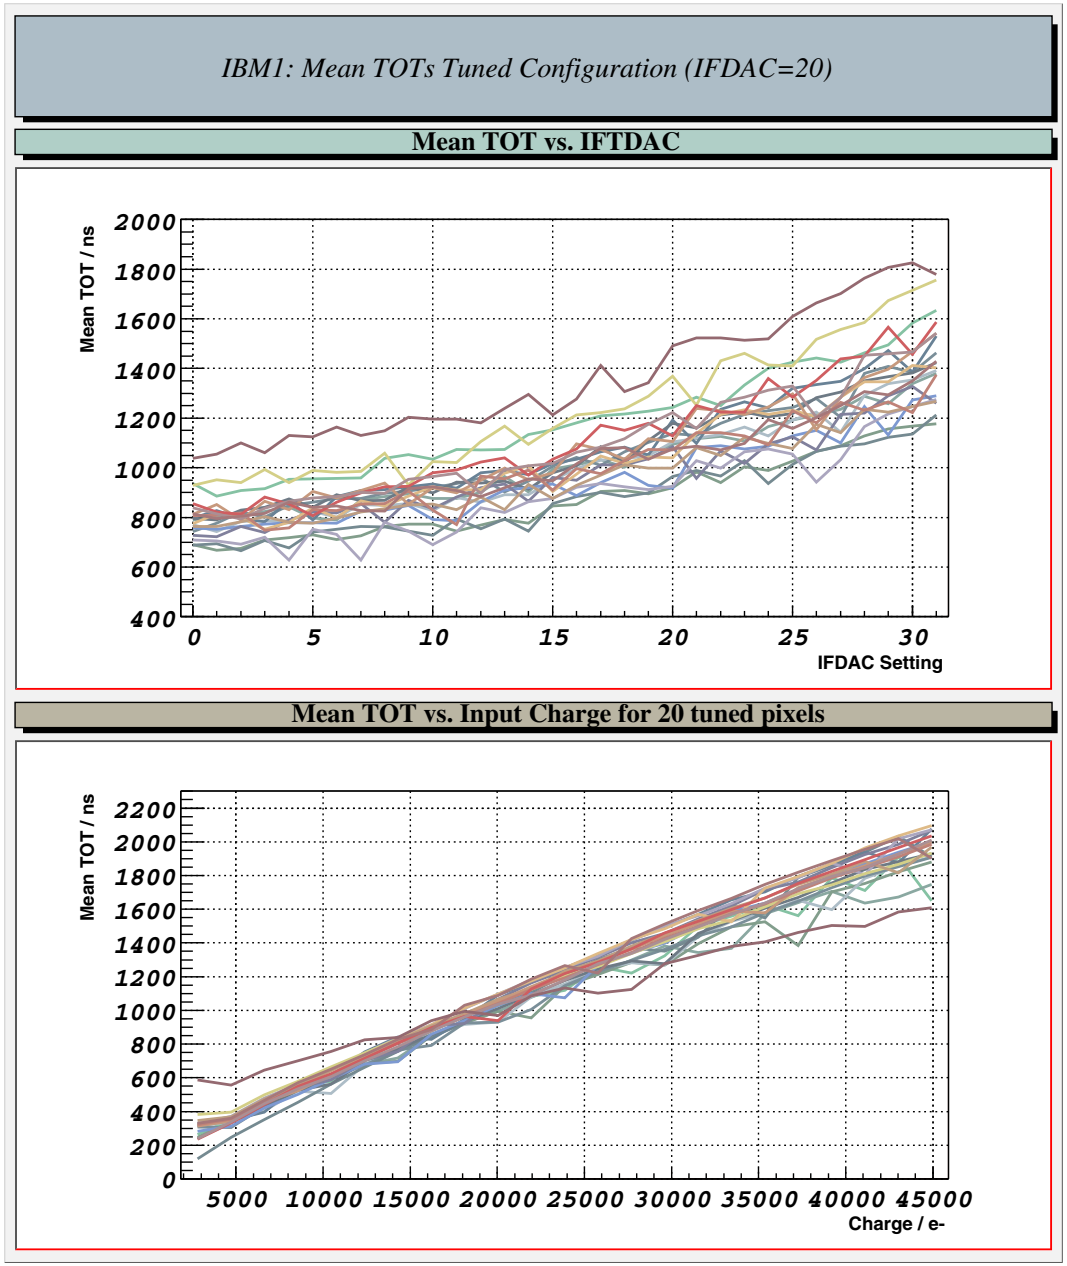

Upper plot shows the TOT versus IFTrim settings. The adjustment is somewhat non-linear, but easily covers the dispersion observed in the testchip.

After tuning, the TOT performance of the different channels is very uniform (in previous pixel chips, we have seen factors of 2 variation in the TOT to charge conversion slope).

- Example of timewalk simulations for IBM with different Cfb values from 5-15fF:

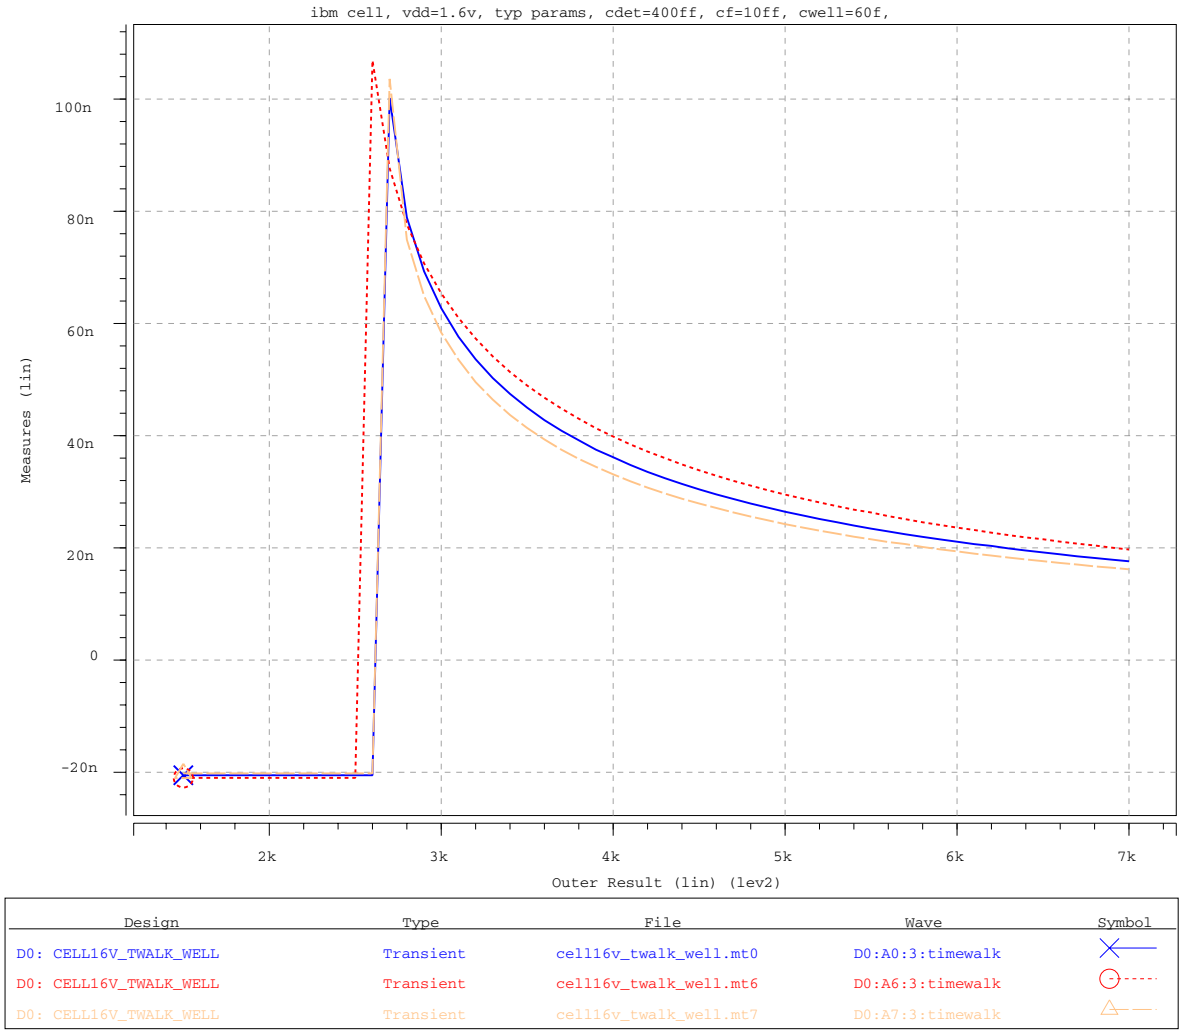

- Timewalk defined relative to 100Ke input. Relevant value is overdrive required to achieve timewalk less than 20ns. Without parasitics, see less than 1000e.
- Here, see roughly 4Ke overdrive required due to poor preamp risetime.

- Example of timewalk simulations for TSMC with different Cfb values from 5-15fF:

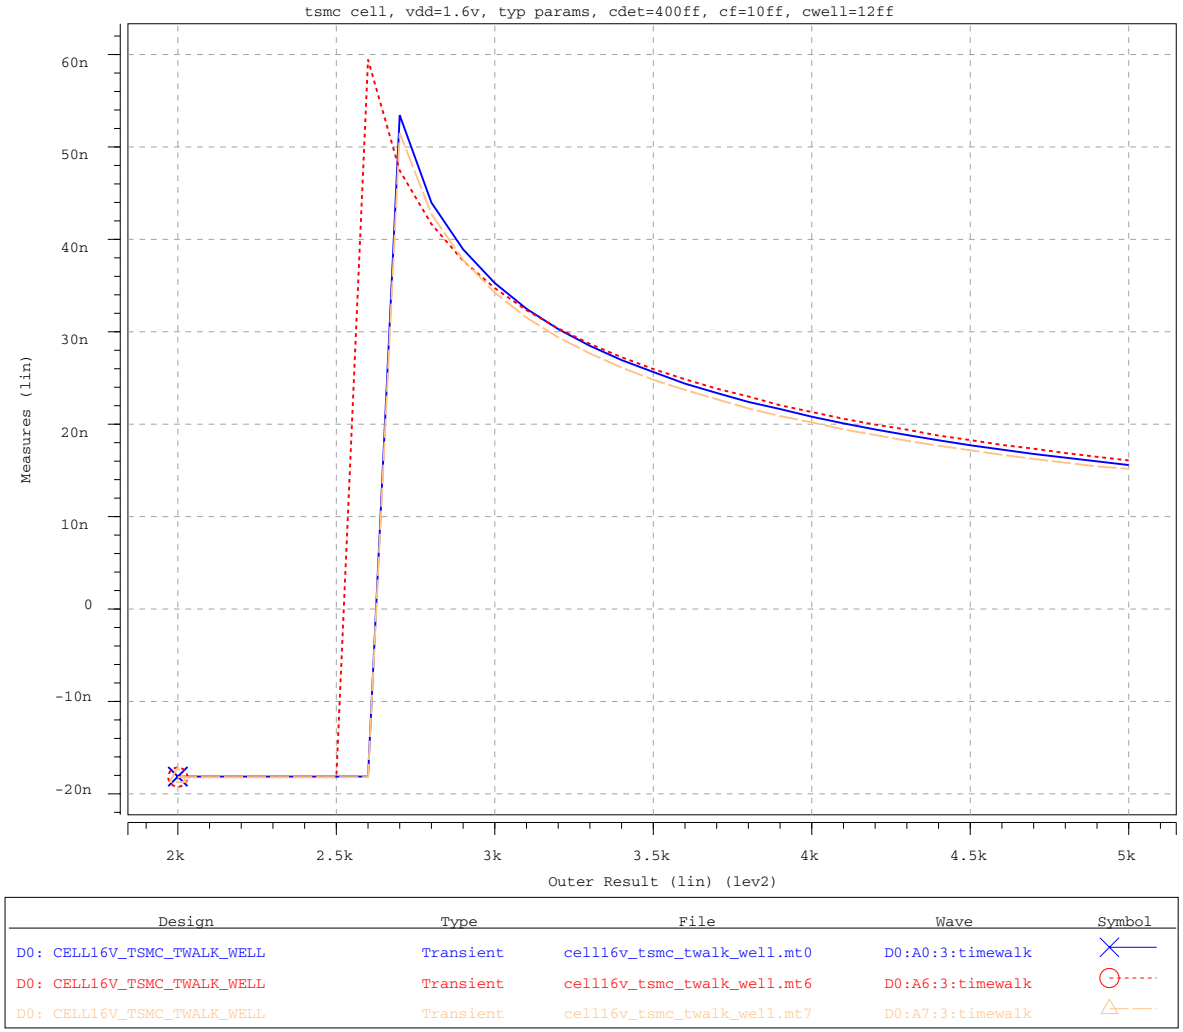

- Here, see roughly 1.5Ke overdrive required. The difference is the well parasitics in IBM (this simulation includes layout parasitics and small TSMC well parasitic). Without parasitics, IBM design is expected to be slightly faster than TSMC.
- Observed value in the lab is about 2Ke for TSMC with this load.

- Example timewalk scan for IBM testchip:

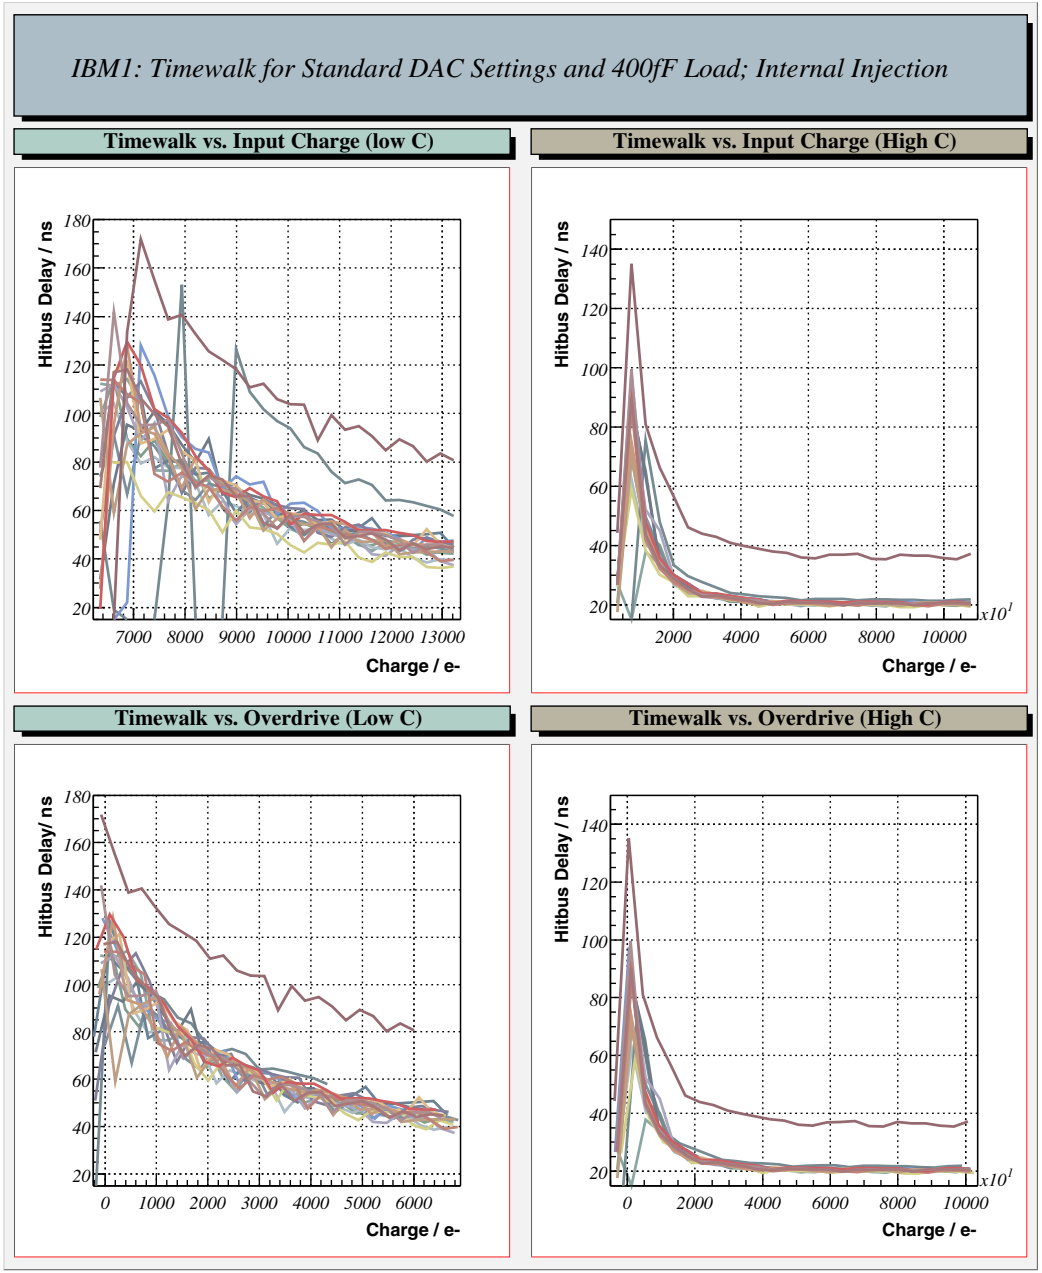

One anomalous channel is the pixel with the 50Ω buffer and test pads loading it.

Other channels all have similar timewalk performance. Measured value is somewhat worse than simulation (measure approximately 6Ke overdrive required for 20ns).

- The major new issue that must be dealt with is the large well-substrate parasitic at the preamplifier output. This degrades the timewalk performance of the IBM version of the front-end to an unacceptable level. However, the TSMC performance is within the acceptable range.
- In order to eliminate this problem, we are presently studying an improved feedback scheme in which only one PMOS is used, and its body is connected to VDDA. This approach appears to lead to identical performance in every respect, with about a 20% loss in dynamic range in the preamp. This will slightly reduce the maximum threshold which we can set. We intend to continue verifying this improved design in simulation over the next 1-2 weeks, and barring unforeseen surprises, we plan to implement it in the final chip.
- One major omission in the test chips was an array of the critical capacitors, which could be used to determine their actual values. IBM provides poor information on the capacitances of their process. We have used analytic calculations to extract a “fringe” capacitance from the layer information, and we have then used this to perform hand calculations of capacitor values. These results do not agree well with the DIVA extracted values. The DIVA extraction seems to provide accurate total capacitances for digital layouts (long thin traces), but does not give accurate values for large analog capacitors (complex geometries with large area and fringe contributions). Our hand calculations suggest that Cfb is about 14fF (we had intended to have 10fF), and Clo is about 4fF. These values were used in our comparisons of simulation and measurement, and produce good agreement.

Improvements for the engineering run:

- Our design and its characterization depends critically on three capacitor values. These are two injection capacitors (Clo and Chi) and the feedback capacitor (Cfb). In the final chip, these capacitors are built up by using all 6 available conductive layers in a complex stack. The Poly layer is used for Chi, and also provides a shield between the input and the substrate during normal operation. M1, M3, and M5 are all connected together as the input node for the preamplifier. M2 contains more of Chi and also Clo. M4 contains more of Chi and also Cfb. We have included a charge-pump capacitor measurement circuit in the bottom of the chip which contains 0, 1, 2, and 4 versions of each of these capacitors, in order to allow us to measure these capacitors chip by chip during the wafer probing. We have adjusted Clo to be 5fF, and Chi to be 40fF. We have a pad with Cfb of 5fF and one with Cfb of 10fF, allowing us to chose one of these two gain configurations for the FE chips.
- Simulations have been performed on the performance of the FE as a function of Cfb, over the range 5-15fF. We have looked at noise, TOT, timewalk, and cross-talk behavior. It appears that the only disadvantages of the higher gain (lower Cfb) is that the cross-talk sensitivity is somewhat increased, and the maximum threshold value is reduced from about 10Ke with Cfb=10fF to about 7Ke with Cfb=5fF. To be more quantitative, for a threshold of 2.5Ke, a charge above 37Ke fires a neighbor pixel if Cfb is 15fF, whereas a neighbor charge of 27Ke is required if Cfb is only 5fF.
- The reduction in threshold dispersion expected from the higher gain is far more important than the modest increase in cross-talk sensitivity seen in simulation.

- We have made a number of minor improvements in the analog blocks to tune their performance, based on the test chip results. These include modifying the TDAC and IF DAC ranges, and implementing an improved charge injection scheme with a better internal chopper, and an independent external injection scheme. We have optimized the sizing of the critical capacitors in the design, as described above.

Digital Issues from the Review

Metastability:

- There are a very limited number of places in the FE chip where asynchronous signals become synchronous, and there are risks of metastability. The only serious concern would be whether a control signal could enter a metastable state which would alter the behavior of the chip for an extended period of time.
- We have carefully evaluated this areas in the design, and believe there are no significant metastability problems.
- The first, and most important, asynchronous signal is the hit in a pixel, which is recorded by an RS FF and then synchronized with the 40MHz crossing clock by recording the 8-bit Grey-coded timestamp value at which it occurs. A metastable state could occur for the latched timestamp data, but this would only corrupt one hit, and could not persist for longer times.
- The second is the sparse scan signal used to signal to the CEU logic that there are hits to be read out in a column. Although this is a control signal, it is sensed by a circuit clocked at 40MHz, and any metastable condition would be forgotten after 25ns.
- A final potential problem is the transfer of the data to the CEU, where the relatively slow risetimes involved could cause a violation of setup times on the CEU latches. Simulations show that we have a large timing margin, even for 3 sigma process variations, so we do not anticipate any problems in this area.

DC Power Consumption:

- Concerns were expressed about excess power consumption associated with the precharge readout of the column. In particular, the differential data readout bus is precharged to the metastable point (VDD/2), and so there is a concern that this could propagate into the subsequent gates, and cause large power consumption by leaving inverters with both transistors turned on.
- The present sense amplifier design is a differential pair in which both the ground and VDD connections are floating until it is enabled. This is followed by an inverter in which the ground connection is floating until the sense amplifier is enabled. Studies of this design indicate that the inverter output will always go to VDD when the sense amplifier is not enabled, so we do not produce metastable output levels.

Readout optimization:

- Possible improvements in the operation of the column readout and sense amplifiers were suggested. In particular, it was suggested to disconnect the data bus lines from the sense amplifier during the sense phase, leaving the signal stored on parasitics at the sense amplifier input (standard technique in commercial SRAMs).
- We estimated that the power savings from this was quite modest. The speed improvement was significant, but simulations already show a factor of two margin in the data transfer speed, so no modification was made. However, an extra pipeline stage was added to the timing of the TOT processor which follows the sense amplifiers, in order to make its timing more conservative.

Issues related to the internal decoupling capacitors

- In order to suppress the large digital transients produced by 0.25µ annular cells (min-size inverter in standard cell library produces 1mA spike when it switches !), we have decided to implement “smart” decoupling capacitors locally in the pixel matrix.
- The basic design is based on the IBM recommendations (PMOS in grounded NWell with a large series NMOS to ground for controlling capacitor). We have added considerable intelligence to the capacitor to deal with potential yield problems (we are running low-power digital chips where 10mA is a significant change in the power budget).
- The basic capacitor cell includes a capacitor with a value of about 1.8pF, a large NMOS, and a thin M1 trace on the minus side. The M1 trace has a resistance of about 20Ω and the NMOS when turned on (Vds=2.0V) has a resistance of about 30Ω, giving an effective 50Ω series resistance with each unit cell. We believe this will mitigate any LC resonant effects due to the inductive bond wires for the power. If the unit cell contained only the capacitor and the NMOS switch, a shorted capacitor could cause up to 10mA of local current flow (Id of NMOS).
- There is a “CapTest” active high signal which puts the capacitor into a test state where it is disconnected from the power rail, and its leakage is compared to a global 15nA bias circuit. If the capacitor leakage is larger than the threshold, then when CapTest returns to inactive, an internal latch keeps this bad capacitor disconnected from the ground net.

- There is logic for an “AutoTest” feature. In this case, a local current in excess of about 1mA during normal operation (CapTest not active) will cause the testing logic to be forced permanently on, resulting in the cell drawing the 15nA test current, but not the large current through the shorted capacitor. This will allow automatic disabling of bad capacitors even without a CapTest cycle being run.
- There is PowerOn circuitry to make sure that CapTest is high during powerup. This will make sure that all capacitors are disconnected from the VDD net when power is first applied. The capacitors will then all charge up to VDD via the small 15nA local test current. Even if all capacitors are defective, this will ensure that they cannot short the global VDD net. Also, the capacitor logic is on a separate VDD net connected at the bottom of chip. This guarantees that internal voltage drops in the chip due to local capacitor shorts cannot reduce the VDD for the control logic to the point where it no longer operates correctly.
- We believe that this combination of safety features should reduce the risk for implementing this new feature to an acceptable level.
- A total of 3 smart capacitors are placed in each pixel, for a total of 8640 in the chip, giving roughly 15nF of total decoupling. These capacitors are claimed by IBM to have excellent properties up into the GHz region.

Layout of pixel, showing two FE blocks and two cap groups:

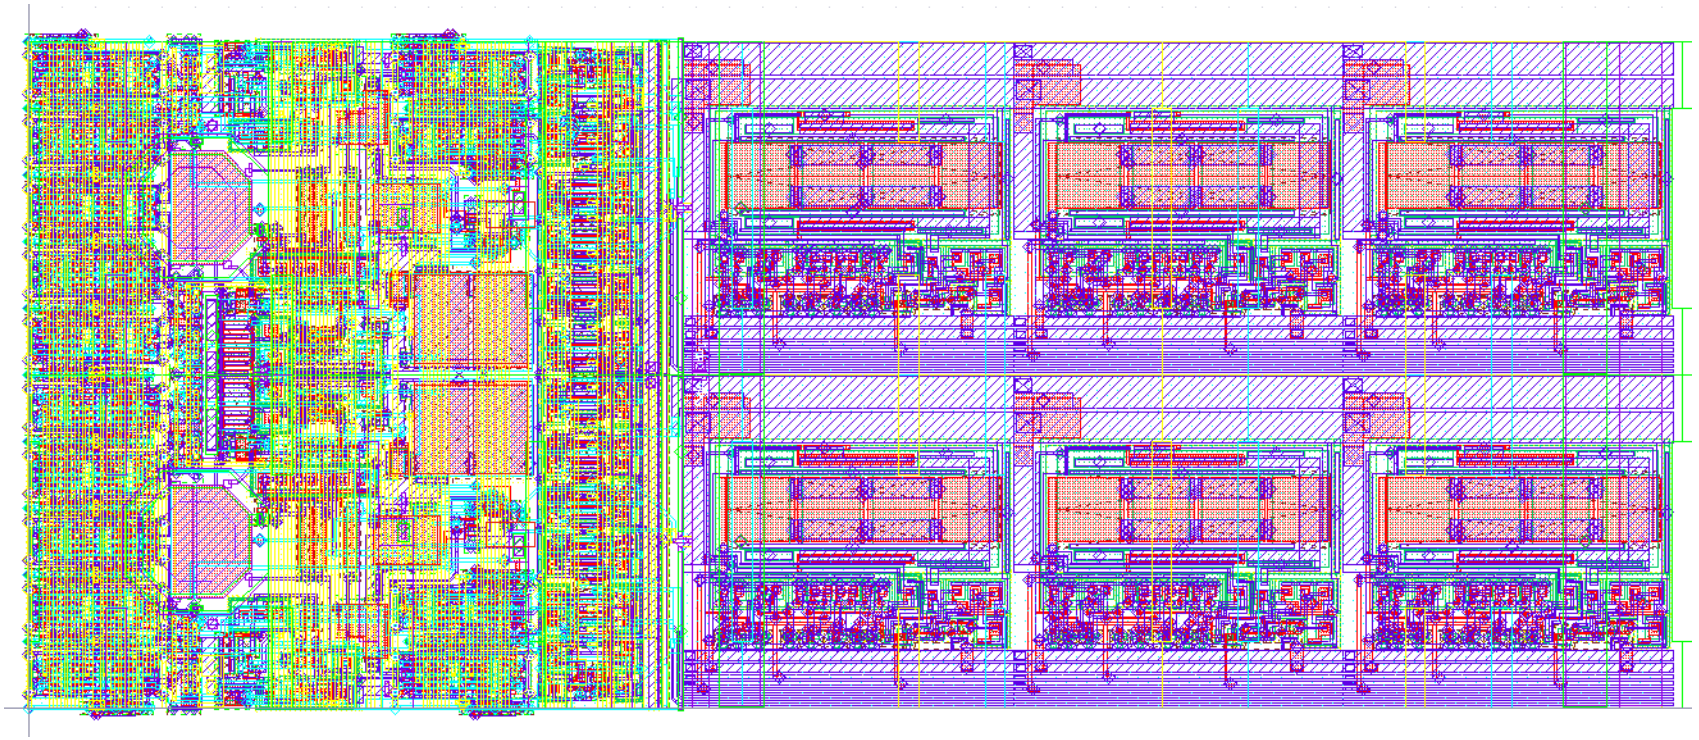

- Capacitor size is roughly 40x50μ, allowing the placement of 3 capacitors in the remaining empty space in the pixel.

Progess in verification of top-level design

Verilog (functional verification):

- Have been running top-level Verilog simulations on digital part of chip, including exercising the Command Decoder, Global Register and Pixel Register.
- Have injected simple hit patterns and seen correct data appear at output.
- Next steps involve creating more sophisticated test vectors for the hit readout, and more complete testing of results for correctness.
- Much careful work done here, and no problems are expected.

TimeMill and PowerMill (timing verification):

- Have been running column pair simulation (without the readout control at the bottom of the chip) for about one week. This is a fully extracted netlist, including about 125K FETs and 250K parasitic capacitors. This seems to operate correctly.
- The power consumption looks a bit higher than expected. PowerMill predictes 32mA DC current for the digital readout in an idle state (no hits being processed), increasing to about 50mA if all column pairs are active, with an average rate of about 1 hit per chip in each 25ns crossing. These numbers require further checking, as occasionally some memory nodes power up into metastable states which draw high current until written into (SRAM in EOC and pixel are not reset in this design). Our nominal power budget was initially set to be 25mA with a worst case of 40mA.

- First simulations including the readout controller have been performed, and correct hits are seen. This is not yet an extracted simulation for the readout controller and its connections.
- Next steps are to create a final netlist for the single column pair with digital controller, including all parasitics. In addition, a final netlist for the “little chip”, which includes all nine column pairs, but only 16 pixels in each, and two EOC buffers in each, will be made. These two netlists should allow us to verify all of the critical timing of the digital readout, and yet be small enough for convenient simulation with TimeMill.
- This work will continue over the next 2-3 weeks until the anticipated submission date.

Present Status of FE-I

Overall Layout:

- All blocks are complete, including special power management prototype blocks.
- Final integration of synthesized digital blocks at the bottom of the chip has just been completed.
- Only block not yet integrated is prototype linear regulator, presently in layout in Bonn.
- LVS and DRC performed so far only on “little chip” consisting of 2 EOC buffers per column pair and 16 pixels per column (about 300K transistors).
- Hercules checks made on other small chips, and pad frame in particular looks OK.

Overall Schematic:

- Top level schematic recently had final major integration occur, to incorporate the synthesized digital readout and command decoder block. It is now essentially complete.
- LVS checking of little chip is ongoing, and for full chip should start soon.

Expect to finish layout and schematic in the next few days...

Brief tour of the layout:

- Top level view of the chip (all 5 metals displayed):

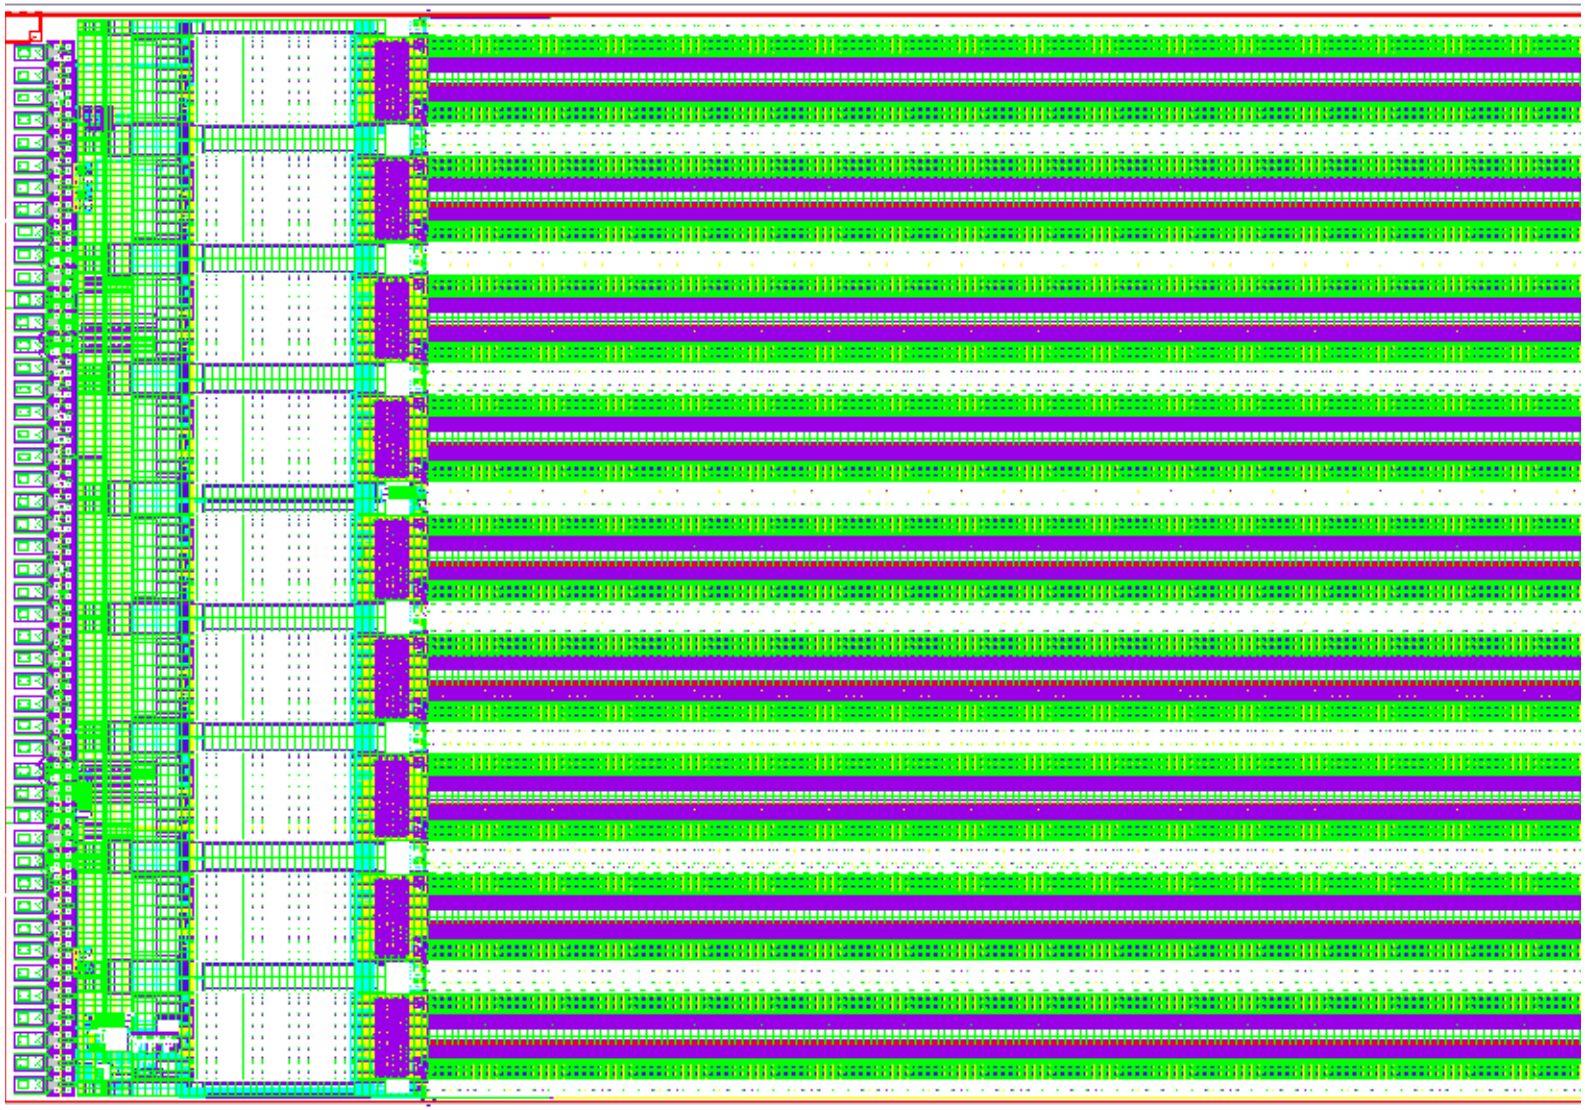

- Zoom into bottom of column region, showing integration of DACs and bias cells with analog columns, and CEU+TOT processor with digital columns:

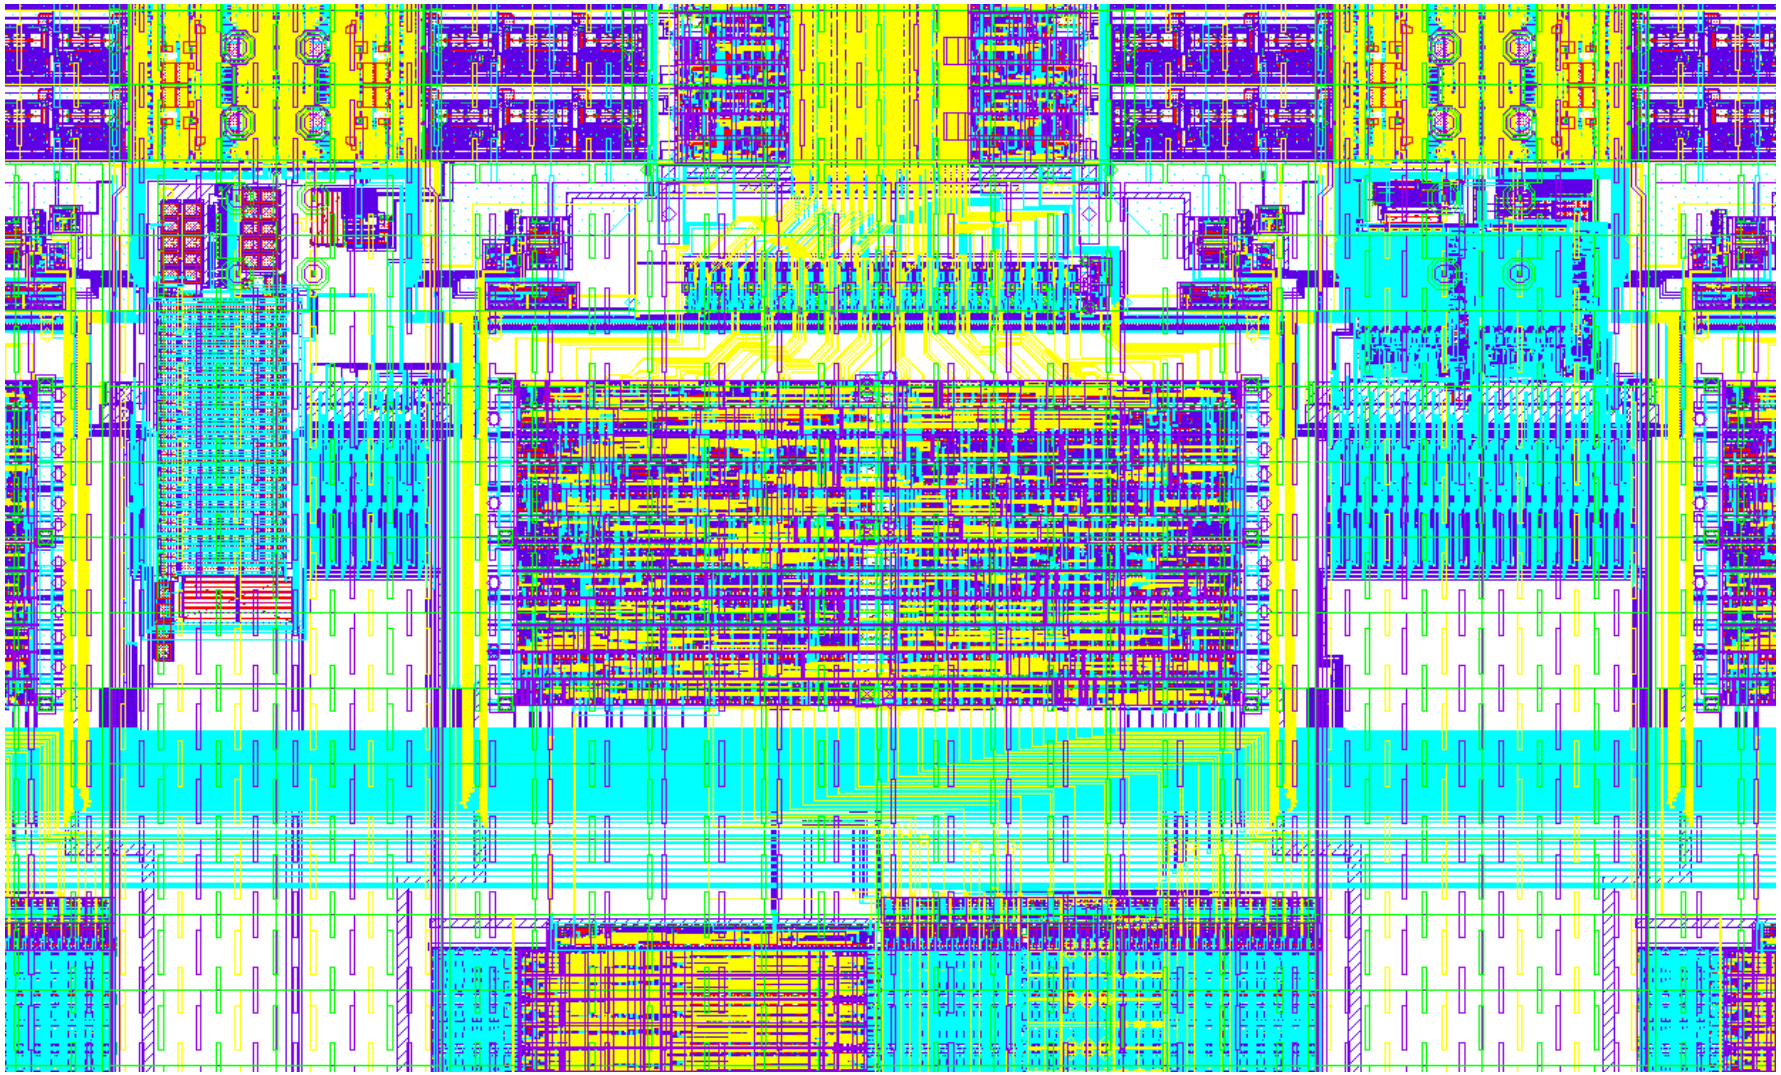

- Left analog column has current reference and register bits, right has pair of 8-bit DACs and register bits (all registers use SEU-tolerant latches).

- Zoom into EOC buffer blocks, each containing 64 hit buffers for a column pair (requiring a total vertical height of about 1mm):

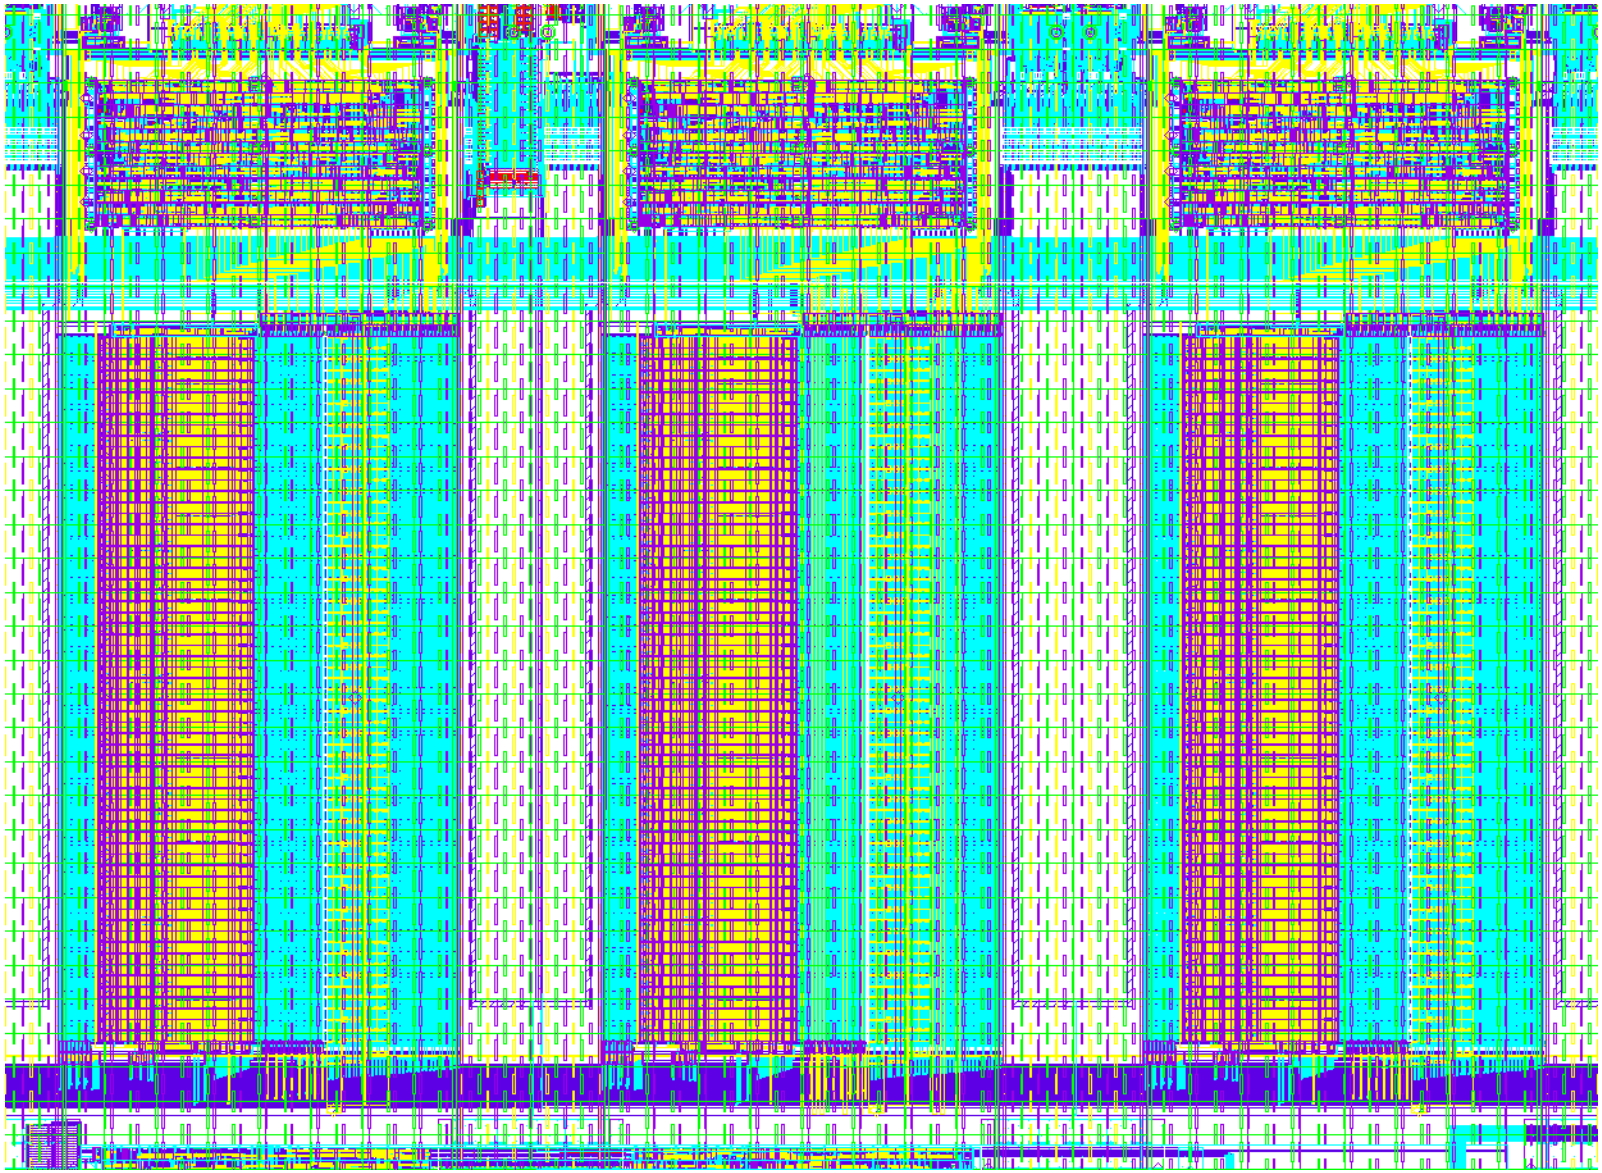

- TOT processor blocks feed into EOC blocks, horizontal bus is at bottom.

- Zoom showing EOC blocks and bottom of chip, including synthesized command decoder and readout controller blocks:

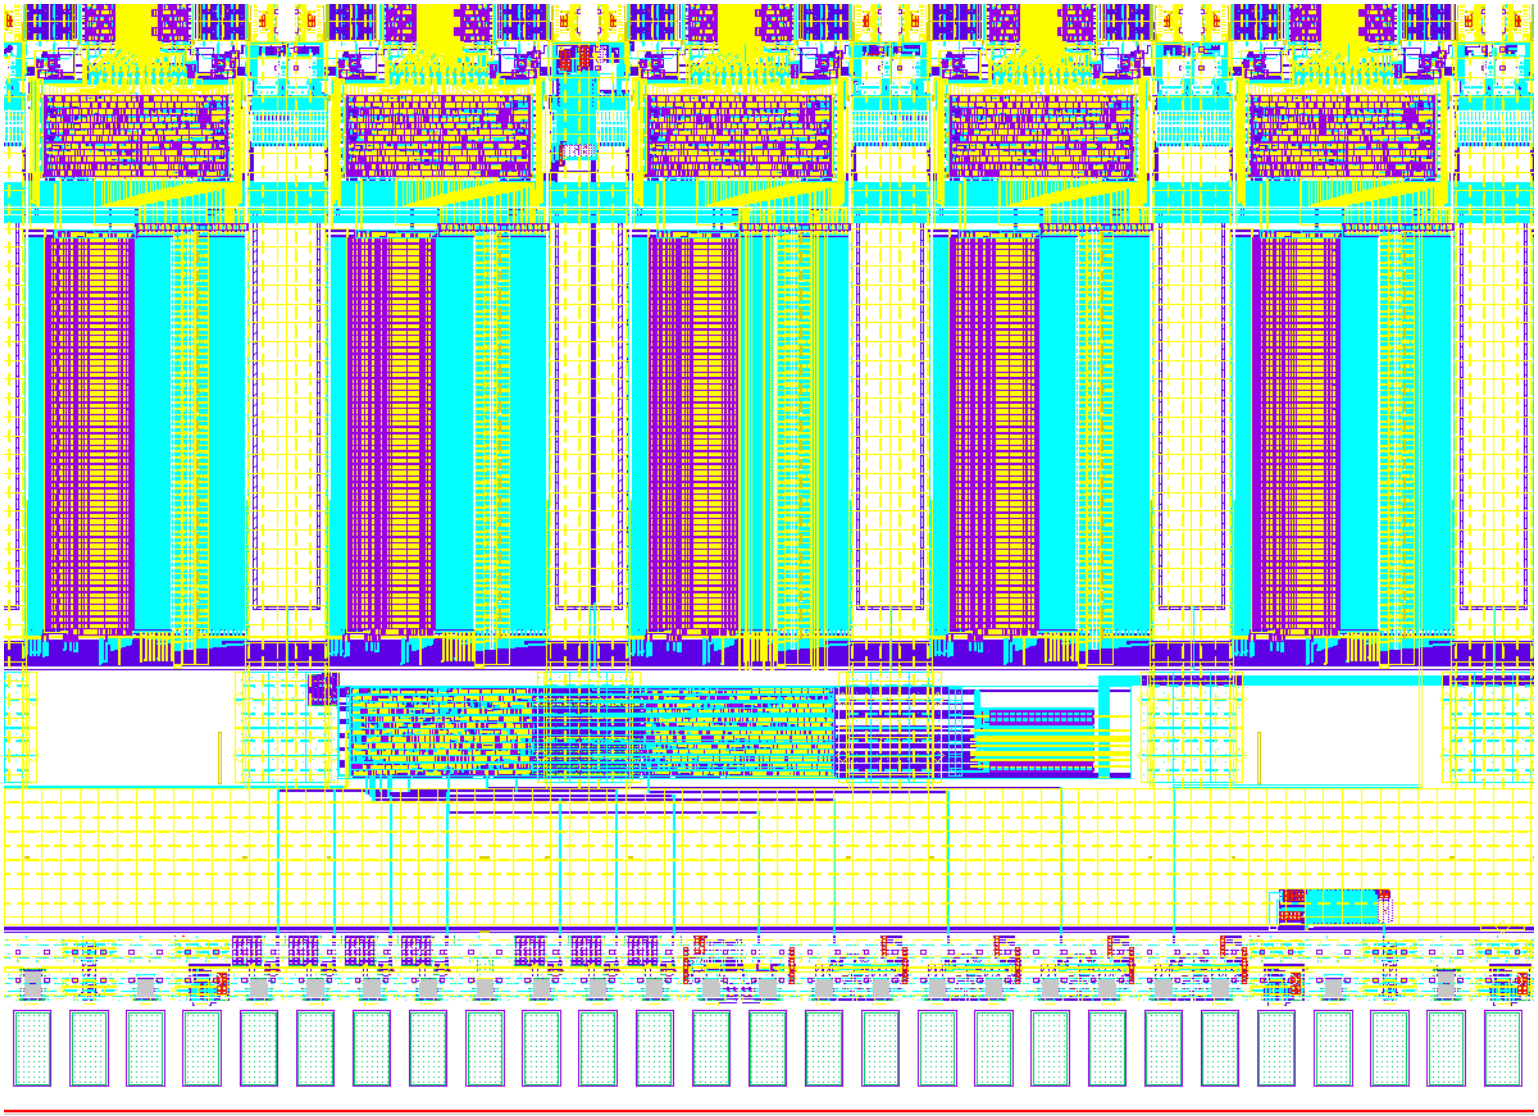

- Note that the bottom of the chip is still largely empty.

- Zoom into Pixel FE block:

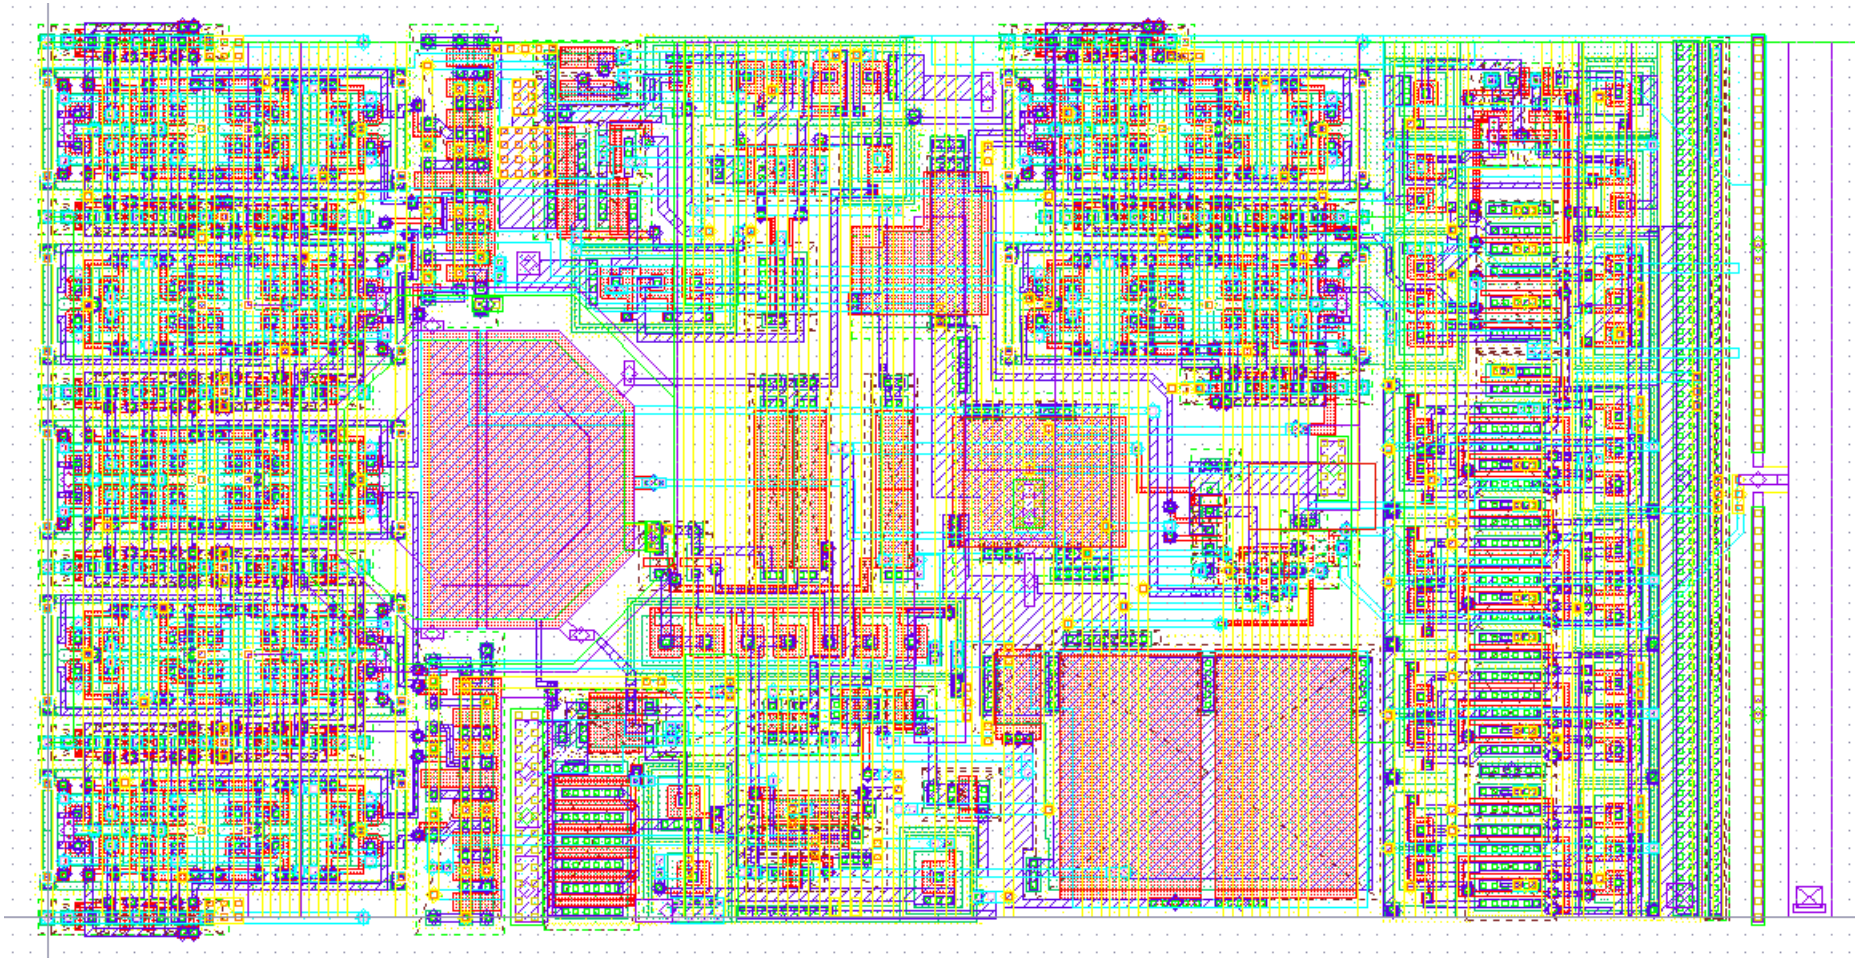

- Left of bump, can see 10 SEU-tolerant latches. Lower right below bump is preamp, center is feedback, top is second stage and discriminator. Right end includes leakage compensation capacitors and additional 4 latches and logic for control of hits, calibrations, and digital injection.

- Zoom into readout region of pixel (two back-to-back columns):

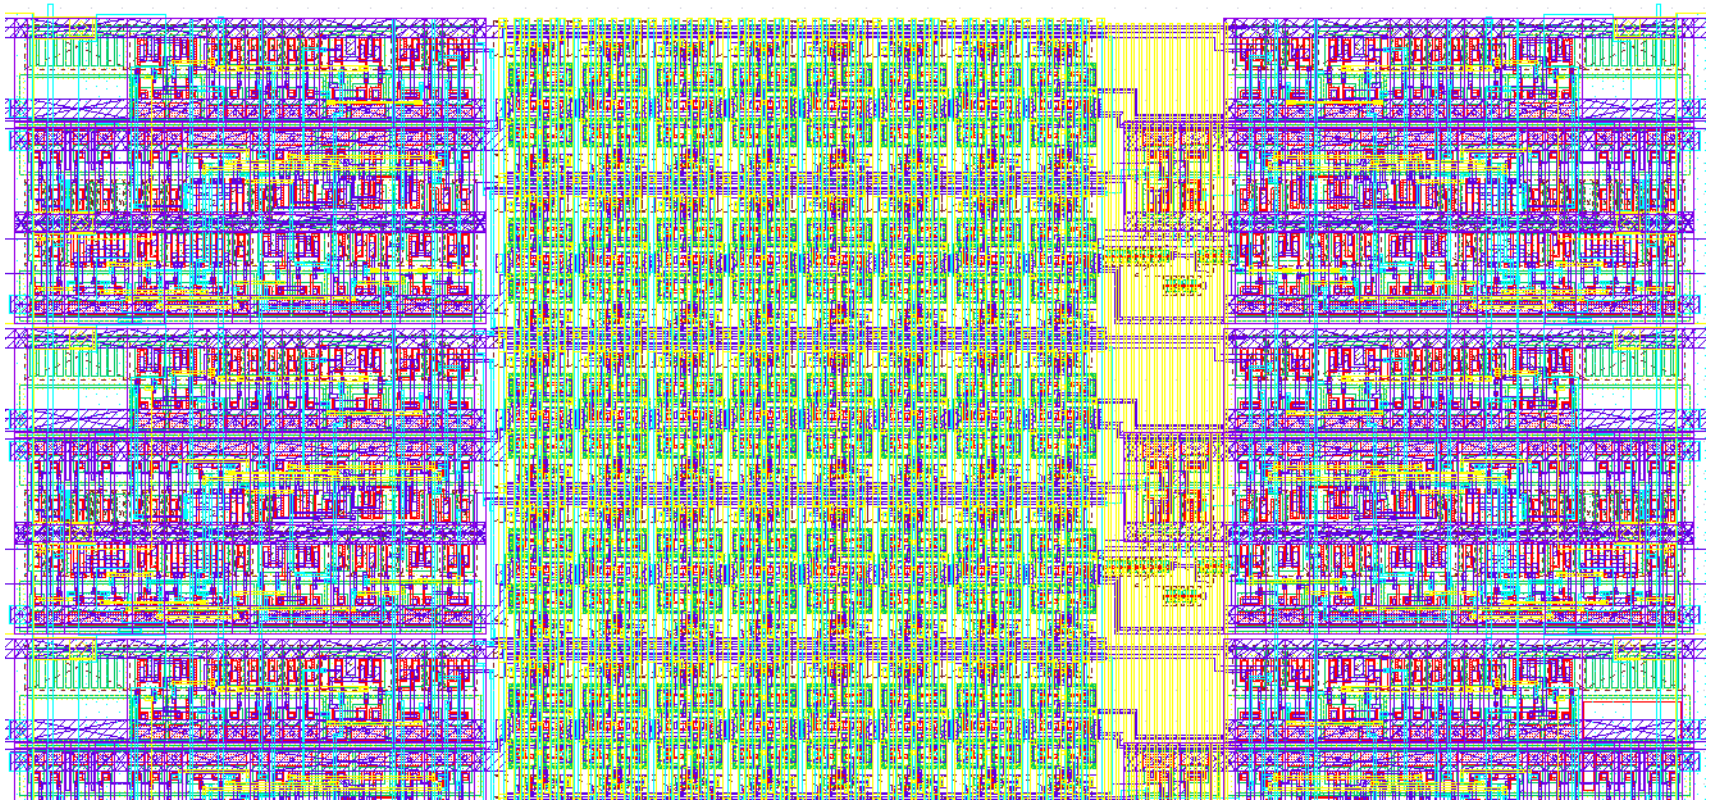

- Central region includes dual 8-bit differential SRAM for LE and TE information for each pixel plus address ROM. Everything is differential (timestamp input, plus RAM and ROM output).
- Left and right sides contain hit logic, sparse scan, and handshaking with CEU for data transfer.

FE-I Pinout and Geometry

Sketch of pin assignments and overall geometry of die:

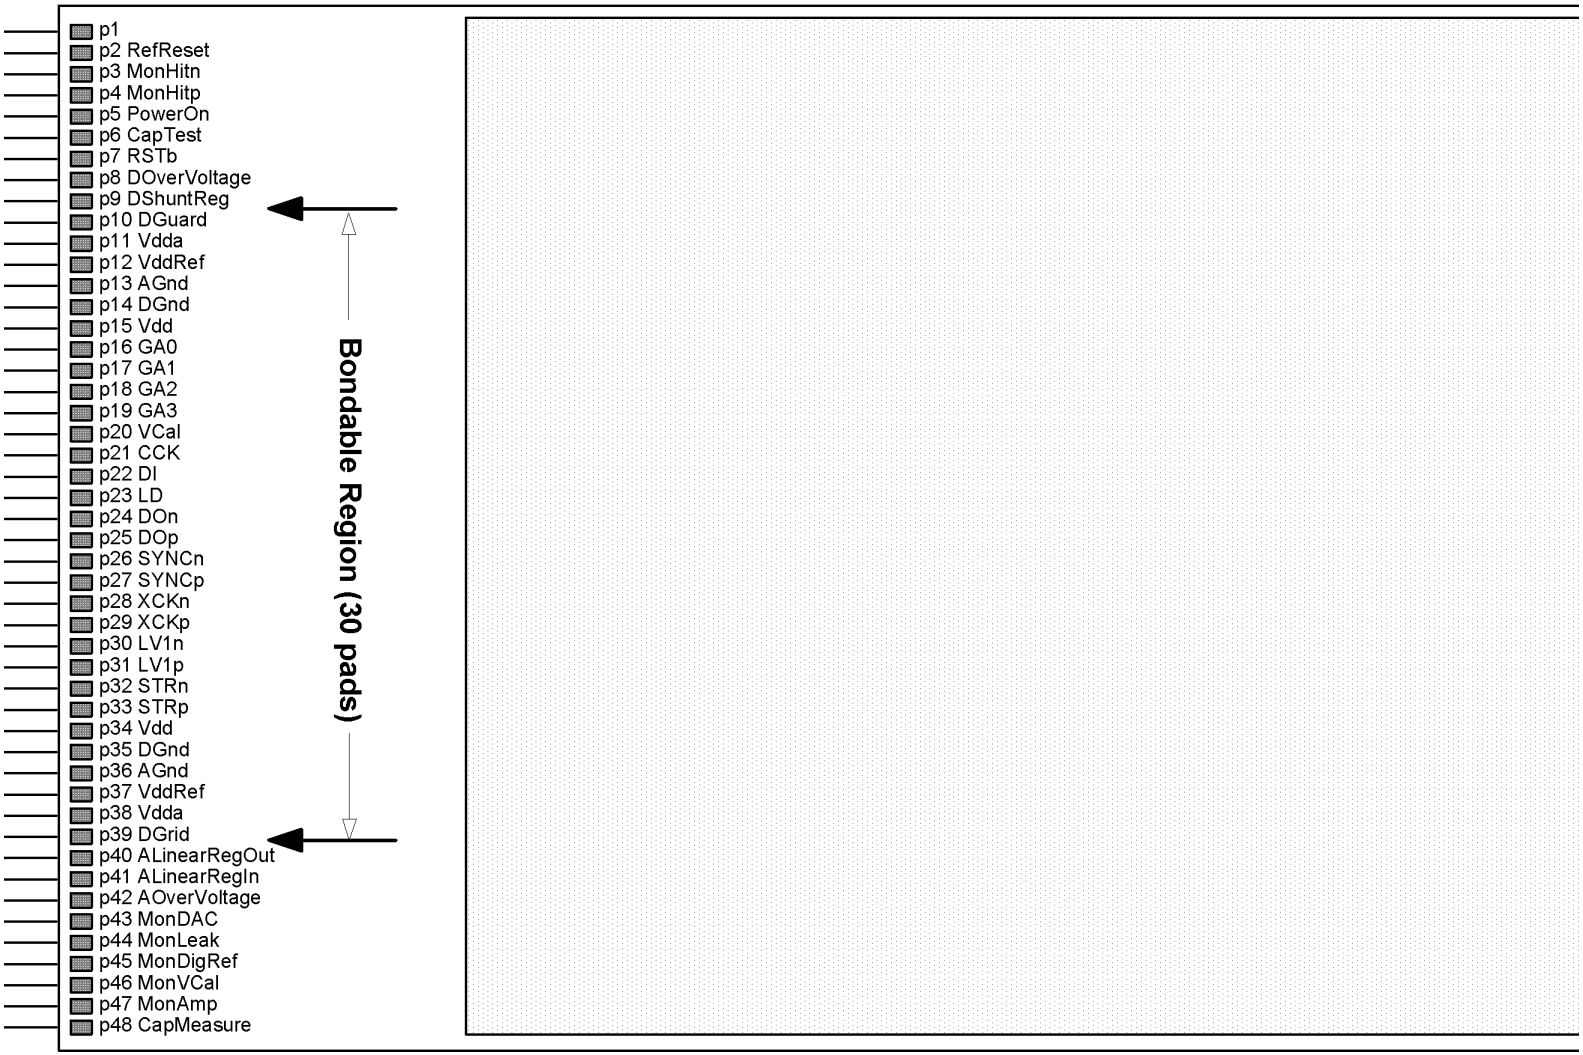

- Note new 100µ x 200µ pad geometry to reduce effects of probing on bonding, and to allow multiple bonding attempts before pad damage becomes too severe.

New features and pins outside the bonding region:

- RefReset pin is active high pin to apply a reset only to the current references. This should never be needed, but is a safety feature.
- CapTest pin is to control new “smart” decoupling capacitors which have been placed internally in the FE chip. Three capacitors, totalling about 5.5pF, are connected to the digital supply, and placed inside each pixel. This provides a total of about 16nF of decoupling in the chip. It should strongly suppress any voltage transients generated by the operation of digital logic in the pixel matrix. The test pin would normally be pulsed in order to cause the capacitors to check for an internal-short condition. Shorted capacitors would be automatically disconnected from the power supply.
- PowerOn pin provides a power-on reset, which can be bonded across to the RSTb pin for testing. The Global Register, when reset, suppresses the basic clocks for the digital readout. The major power consumption in the digital logic is in the TSI distribution, CEU operation, and EOC buffer state machines. When the Global Register is reset, the digital part of the chip will operate in a low-power condition (less than 10% of nominal power), but all basic registers will still operate. The analog supply could be turned on, but since the DAC values would all be set to 0, the current consumption would also be very low. This low-power power-on state would permit simple continuity tests of the module (and the rest of the cable plant) with a power consumption so low that no cooling would be required.

- Power Management features: Two overvoltage clamp circuits are included, one for each power supply. They use a diode and a resistor to set a soft threshold of about 2.7V, after which a large PFET is used to sink excess voltage to ground. Note that the power pads also include the recommended IBM transient clamps, which are designed to protect the chip from sharp spike transients on the power rails either with or without power applied to the chip. In addition to the overvoltage clamps, there are two simple regulators. One is a shunt regulator, based on the same circuit as the clamp, but with a threshold of 2.0V. The second is a simple linear regulator using a band-gap reference, and set for 1.6V operation. The shunt regulator is intended for study of powering schemes based on constant current supplies. The linear regulator would generate the analog supply voltage from the digital supply voltage, allowing operation of the FE chip on a single power supply. There is little risk posed by these circuits if the wire-bonds are not connected, and placing them inside the FE chips allows the performance of modules to be compared with and without these circuits, without changes to the Flex design.
- MonDAC provides multiplexed access to all of the internal DACs for characterization during testing.
- MonLeak provides access to a current summing tree (controlled in the same way as the HitBus) that allows a direct measurement of the preamp feedback current and the sensor leakage current: I(OutLeak) = 3*If + ILeak. This has already proven very useful in chip characterization. A simple internal ADC, based on a 9-bit DAC, is also provided.

- MonRef allows direct monitoring of the current reference used for the LVDS I/O pads, without requiring any other circuitry to operate on the FE chip.
- MonVCal allows direct monitoring of the VCal voltage generated internally on the FE chip for charge injection calibrations. VCal is generated by a 9-bit current DAC and a resistor. The resistor is matched to the one used in the current reference, providing first-order cancellation of process variations.
- MonAmp would be upgraded to allow us to see the preamp waveform, the two sides of the second amplifier, and the chopper input. There is also a 100Ω buffer amplifier, which could drive a daisy-chained bus of test amplifiers, provided only one was enabled at any time. This circuitry has proven vital in Analog Test chip.
- CapMeasure pin is attached to new capacitor measurement circuitry, which uses a charge pump circuit to measure accurate values for the critical capacitors used in the front-end (C(feedback), C(inj-low), C(inj-high)) by measuring a single DC current. This circuit has been used in the DMILL CapTest chip, and can provide accurate measurements of capacitor arrays at the fF level.

- Lower right chip corner including several analog blocks (50W buffer, CapMeasure circuit, MonLeak ADC, LVDS current reference, and overvoltage protection). Will also include linear regulator block:

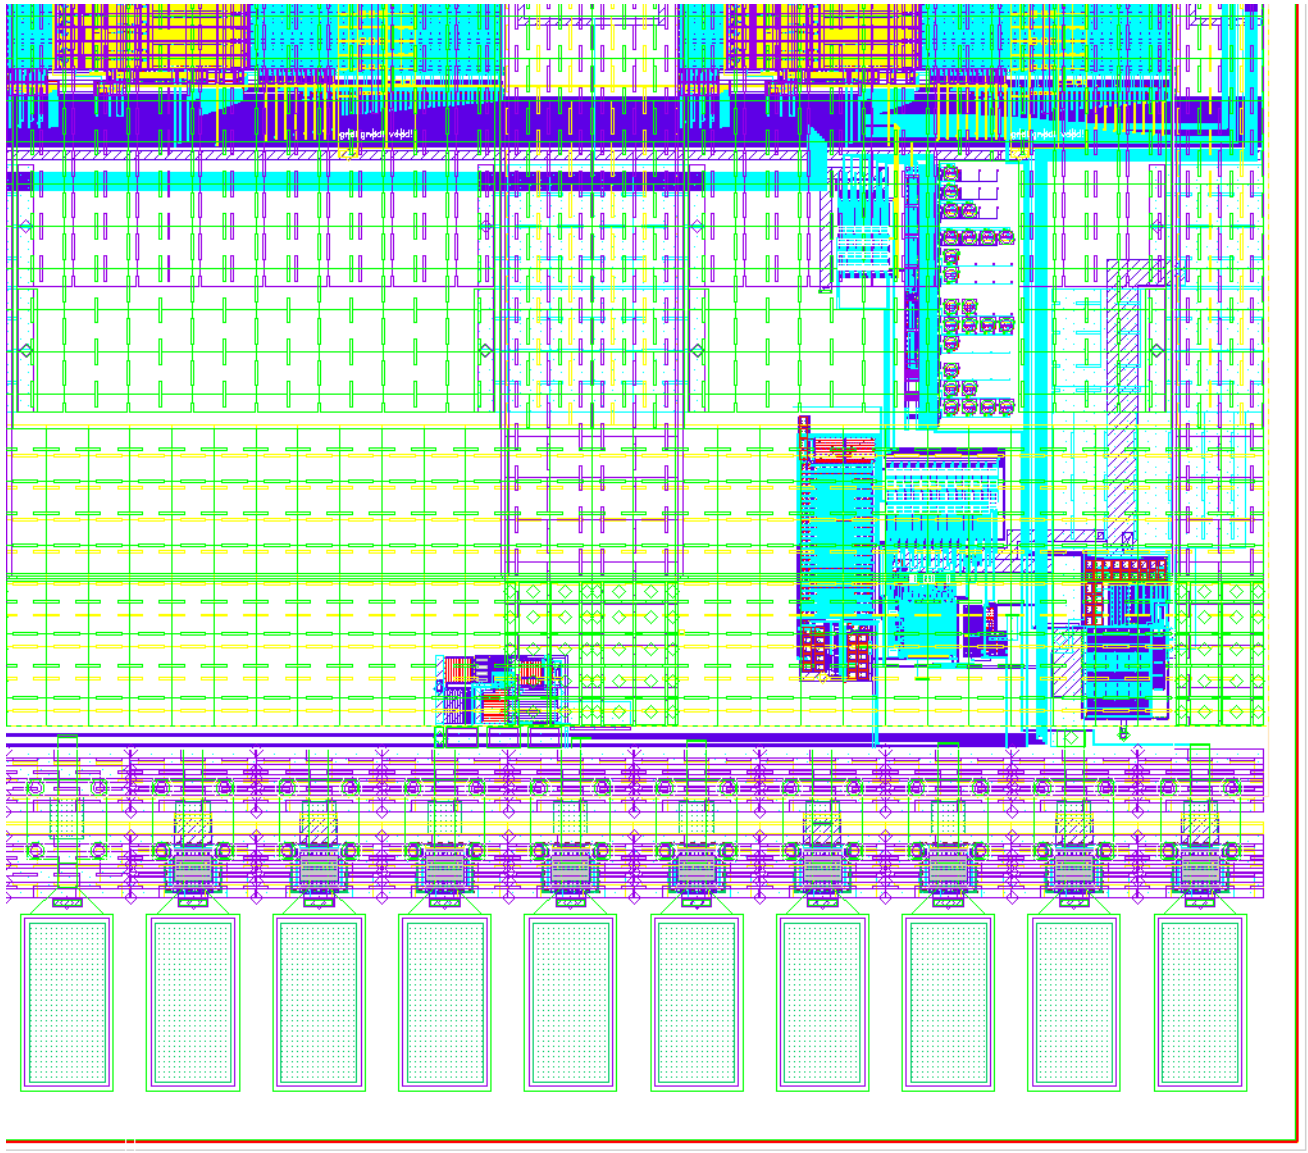

Reticle for FE-I Run

Similar to FE-D2 run, all designs should use same pads

- **Two FE-I chips:** the plan is that these would have slightly different FE designs.
- **MCC-I chip:** this is the complete new MCC with U-pinout to satisfy module constraints. It satisfies all of our presently known production requirements.
- **DORIC-I and VDC-I chips:** they are improved versions of the designs submitted in the Feb MPW run, including 4-channel VDC matched to Taiwan opto-package.
- **Analog Test Chip:** this is very similar to the test chips fabricated in Feb/Mar with IBM and TSMC, but contains the final design and layout of all analog blocks, and 64 pixels instead of 20. We are attempting to keep a similar pinout.
- **LVDS Buffer Chip:** this is a convenient way to include the interface between a single chip and our test system into a rad-hard chip. Given the absence of commercial LVDS drivers operating at less than 3.3V, this is essential. It contains 4 LVDS->CMOS converters, 3 LVDS->LVDS repeaters (3.5mA outputs).
- **PM bar:** may be useful for checking details of device characteristics, although the very good parameter stability seen so far suggests it may no longer be needed. We plan to include the small bar designed by the CERN group, and used by them to track the parameter stability on all multi-project runs. We have already used such bars from our MPW run to characterize transistors and check the IBM SPICE models, and irradiations have been performed to study the process as well.

Preliminary reticle layout:

IBM Reticle Layout (drawn to scale, 100u gap between Chipedge of designs)

Nominal Rules:

* ChipEdge is 10u outside ChipGuard, and defines design size given below.

* Different designs are separated by 100μ gap between ChipEdges.

Reticle size is: 14.7 (W) x 14.98 (L) mm.

IBM adds 138μ in one direction and 378μ in the other direction.
We choose to have the 138μ added to the left and right of the reticle shown here.

Reticle stepping increments with these rules are: 14.838 (W) x 15.358 (L) mm.